

Funnel Libraries for Real-Time Robust Feedback Motion Planning

Anirudha Majumdar and Russ Tedrake
Computer Science and Artificial Intelligence Lab
Massachusetts Institute of Technology
Cambridge, MA, USA
Email: {anirudha,russt}@mit.edu

July 21, 2016

Abstract


We consider the problem of generating motion plans for a robot that are guaranteed to succeed despite uncertainty in the environment, parametric model uncertainty, and disturbances. Furthermore, we consider scenarios where these plans must be generated in real-time, because constraints such as obstacles in the environment may not be known until they are perceived (with a noisy sensor) at runtime. Our approach is to pre-compute a library of “funnels” along different maneuvers of the system that the state is guaranteed to remain within (despite bounded disturbances) when the feedback controller corresponding to the maneuver is executed. We leverage powerful computational machinery from convex optimization (*sums-of-squares programming* in particular) to compute these funnels. The resulting *funnel library* is then used to sequentially compose motion plans at runtime while ensuring the safety of the robot. A major advantage of the work presented here is that by explicitly taking into account the effect of uncertainty, the robot can evaluate motion plans based on how vulnerable they are to disturbances.

We demonstrate and validate our method using extensive hardware experiments on a small fixed-wing airplane avoiding obstacles at high speed (~ 12 mph), along with thorough simulation experiments of ground vehicle and quadrotor models navigating through cluttered environments. To our knowledge, the resulting hardware demonstrations on a fixed-wing airplane constitute one of the first examples of provably safe and robust control for robotic systems with complex nonlinear dynamics that need to plan in realtime in environments with complex geometric constraints.

1 Introduction

Imagine an unmanned aerial vehicle (UAV) flying at high speed through a cluttered environment in the presence of wind gusts, a legged robot traversing rough terrain, or a mobile robot grasping and manipulating previously unlocalized objects in the environment. These applications demand that the robot move through (and in certain cases interact with) its environment with a very high degree of agility while still being in close proximity to obstacles. Such systems today lack guarantees on their safety and can fail dramatically in the face of uncertainty in their environment and dynamics.

The tasks mentioned above are characterized by three main challenges. First, the dynamics of the system are nonlinear, underactuated, and subject to constraints on the input (e.g. torque limits). Second, there is a significant amount of uncertainty in the dynamics of the system due



(a) A plane deviating from its nominal planned trajectory due to a heavy cross-wind.

(b) The “funnel” of possible trajectories.

Figure 1: Not accounting for uncertainty while planning motions can lead to disastrous consequences.

to disturbances and modeling error. Finally, the geometry of the environment that the robot is operating in is unknown until runtime, thus forcing the robot to plan in *real-time*.

1.1 Contributions

In this paper, we address these challenges by combining approaches from motion planning, feedback control, and tools from Lyapunov theory and convex optimization in order to perform robust real-time motion planning in the face of uncertainty. In particular, in an offline computation stage, we first design a finite library of open loop trajectories. For each trajectory in this library, we use tools from convex optimization (*sums-of-squares (SOS) programming* in particular) to design a controller that explicitly attempts to minimize the size of the worst case reachable set of the system given a description of the uncertainty in the dynamics and bounded external disturbances. This control design procedure yields an outer approximation of the reachable set, which can be visualized as a “funnel” around the trajectory, that the closed-loop system is guaranteed to remain within. A cartoon of such a funnel is shown in Figure 1(b). Finally, we provide a way of *sequentially composing* these robust motion plans at runtime in order to operate in a provably safe manner in previously unseen environments.

One of the most important advantages that our approach affords us is the ability to choose between the motion primitives in our library in a way that takes into account the dynamic effects of uncertainty. Imagine a UAV flying through a forest that has to choose between two motion primitives: a highly dynamic roll maneuver that avoids the trees in front of the UAV by a large

margin or a maneuver that involves flying straight while avoiding the trees only by a small distance. An approach that neglects the effects of disturbances and uncertainty may prefer the former maneuver since it avoids the trees by a large margin and is therefore “safer”. However, a more careful consideration of the two maneuvers could lead to a different conclusion: the dynamic roll maneuver is far more susceptible to wind gusts and perturbations to the initial state than the second one. Thus, it may in fact be more robust to execute the second motion primitive. Further, it may be possible that neither maneuver is guaranteed to succeed and it is safer to abort the mission and simply transition to a hover mode. Our approach allows robots to make these critical decisions, which are essential if robots are to move out of labs and operate in real-world environments.

We demonstrate and validate our approach using thorough simulation experiments of ground vehicle and quadrotor models navigating through cluttered environments, along with extensive hardware experiments on a small fixed-wing airplane avoiding obstacles at high speed (~ 12 miles per hour) in an indoor motion capture environment. To the best of our knowledge, the resulting hardware demonstrations on a fixed-wing airplane constitute one of the first examples of provably safe and robust control for robotic systems with complex nonlinear dynamics that need to plan in real-time in cluttered environments.

1.2 Outline

The outline of the paper is as follows. Section 2 discusses prior work; Section 3 provides a brief background on semidefinite and sums-of-squares (SOS) programming, which are used heavily throughout the paper; Section 4 shows how to use SOS programming to compute funnels; Section 5 introduces the notion of a funnel library; Section 6 describes our algorithm for using funnel libraries for real-time robust planning in environments that have not been seen by the robot before; Section 7 presents extensive simulation results on a ground vehicle model and compares our approach with an approach based on trajectory libraries; Section 7 also considers a quadrotor model and shows how one can use our approach to guarantee collision-free flight in certain environments; Section 8 presents hardware experiments on a small fixed-wing airplane in order to demonstrate and validate our approach; Section 9 concludes the paper with a discussion of challenges and open problems.

2 Relevant Work

2.1 Motion Planning

Motion planning has been the subject of significant research in the last few decades and has enjoyed a large degree of success in recent years. Planning algorithms like the Rapidly-exploring Randomized Tree (RRT) [Kuffner and Lavalle, 2000] (along with variants that attempt to find optimal motion plans [Karaman and Frazzoli, 2011] [Kobilarov, 2012]) and related trajectory library approaches [Liu and Atkeson, 2009] [Frazzoli et al., 2005] [Stolle and Atkeson, 2006] can handle large state-space dimensions and complex differential constraints. These algorithms have been successfully demonstrated on a wide variety of hardware platforms [Kuwata et al., 2008] [Shkolnik, 2010] [Sermanet et al., 2008] [Satzinger et al., 2016]. However, an important challenge is their inability to explicitly reason about uncertainty and feedback. Modeling errors, state uncertainty and disturbances can lead to failure if the system deviates from the planned nominal trajectories.

The motion planning aspect of our approach draws inspiration from the vast body of work that is focused on computing motion primitives in the form of trajectory libraries. For example,

trajectory libraries have been used in diverse applications such as humanoid balance control [Liu and Atkeson, 2009], autonomous ground vehicle navigation [Sermanet et al., 2008], grasping [Berenson et al., 2007] [Dey et al., 2011], and UAV navigation [Dey et al., 2014] [Barry, 2016]. The Maneuver Automaton [Frazzoli et al., 2005] attempts to capture the formal properties of trajectory libraries as a hybrid automaton, thus providing a unifying theoretical framework. Maneuver Automata have also been used for real-time motion planning with static and dynamic obstacles [Frazzoli et al., 2002]. Further theoretical investigations have focused on the offline generation of diverse but sparse trajectories that ensure the robot’s ability to perform the necessary tasks online in an efficient manner [Green and Kelly, 2007]. More recently, tools from sub-modular sequence optimization have been leveraged in the optimization of the sequence and content of trajectories evaluated online [Dey et al., 2011, Dey, 2015]. Prior work has also been aimed at *learning* maneuver libraries in the form of “skill trees” from expert demonstrations [Konidaris et al., 2011].

Robust motion planning has also been a very active area of research in the robotics community. Early work [Brooks, 1982] [Lozano-Perez et al., 1984] [Jacobs and Canny, 1990] [Latombe et al., 1991] focused on settings where the dynamics of the system are not dominant and one can treat the system as a kinematic one. The problem is then one of planning paths through configuration space that are robust to uncertainty in the motion of the robot and geometry of the workspace. Our work with funnels takes inspiration from the early work on fine-motion planning, where the notions of *funnels* [Mason, 1985] and *preimage backchaining* [Lozano-Perez et al., 1984] (also known as *goal regression* or *sequential composition*) were first introduced. The theme of robust kinematic motion planning has persisted in recent work [Missiuro and Roy, 2006] [Guibas et al., 2009] [Malone et al., 2013], which deals with uncertainty in the geometry of the environment and obstacles.

2.2 Planning under Uncertainty

In settings where the dynamics of the system must be taken into account (e.g., for underactuated systems), the work on *planning under uncertainty* attempts to address the challenges associated with uncertainty in the dynamics, state, and geometry of the environment. In particular, chance-constrained programming [Charnes and Cooper, 1959] provides an elegant mathematical framework for reasoning about stochastic uncertainty in the dynamics, initial conditions, and obstacle geometry [Blackmore et al., 2006] [Oldewurtel et al., 2008] [Toit and Burdick, 2011] and allows one to generate motion plans with bounds on the probability of collision with obstacles. In settings where the state of the system is partially observable and there is significant uncertainty in the observations, one can extend this framework to plan in the *belief space* of the system (i.e., the space of distributions over states) [Vitus and Tomlin, 2011] [Vitus and Tomlin, 2012]. While these approaches allow one to explicitly reason about uncertainty in the system, they are typically restricted to handling linear dynamical systems with Gaussian uncertainty. This is due to the computational burden of solving chance constrained problems for nonlinear and non-Gaussian systems.

Similarly, other approaches for belief space planning [Bry and Roy, 2011] [Platt et al., 2010] [Prentice and Roy, 2007] [Berg et al., 2012] [Agha-Mohammadi et al., 2014] [Patil et al., 2015] also approximate the belief state as a Gaussian distribution over state space (however, see [Platt et al., 2012] for an exception to this) for computational efficiency and hence the true belief state is not tracked. Thus, in general, one does not have robustness guarantees. The approach we take in this work is to assume that disturbances/uncertainty are *bounded* and provide explicit bounds on the reachable set to facilitate safe operation of the *nonlinear* system.

More generally, the rich literature on Partially Observable Markov Decision Processes (POMDPs)

[Kaelbling et al., 1998] is also relevant here. POMDPs present an elegant mathematical framework for reasoning about uncertainty in state and dynamics. However, we note that our focus in this work is on dynamical systems with continuous state and action spaces, whereas the POMDP literature typically focuses on discretized state/action spaces for the most part.

2.3 Reachable Sets

Reachability analysis for linear and nonlinear control systems has a long history in the controls community. For *linear* systems subject to bounded disturbances, there exist a number of techniques for efficiently computing (approximations of) backwards and forwards reachable sets [Kurzhanski and Varaiya, 2000] [Girard, 2005] [Yazarel and Pappas, 2004]. One can apply techniques from linear reachability analysis to conservatively approximate reachable sets of nonlinear systems by treating nonlinearities as bounded disturbances [Althoff et al., 2008]. This idea has been used in [Althoff and Dolan, 2014] for performing online safety verification for ground vehicles. A similar idea was used in [Althoff et al., 2015] to perform online safety verification for UAVs and to decide when the UAV should switch to an emergency maneuver (a “loiter circle”). In this paper we will also compute outer approximations of reachable sets (which we refer to as “funnels”). However, the approach we present here is not based on a linearization of the system and thus has the potential to be less conservative for highly nonlinear systems. Further, the scope of our work extends beyond verification; the emphasis here is on safe real-time *planning* with funnels.

Approximations of reachable sets for nonlinear systems can be computed via Hamilton-Jacobi-Bellman (HJB) differential game formulations [Mitchell et al., 2005]. This method was used in [Gillula et al., 2010] for designing motion primitives for making a quadrotor perform an autonomous backflip in a 2D plane. While this approach handles unsafe sets that the system is not allowed to enter, it is assumed that these sets are specified *a priori*. In this paper, we are concerned with scenarios in which unsafe sets (such as obstacles) are not specified until runtime and must thus be reasoned about *online*. Further, techniques for computing reachable sets based on the HJB equation have historically suffered from the curse of dimensionality since they rely on discretizations of the state space of the system. Hence, unless the system under consideration has special structure (e.g., decoupled systems [Chen and Tomlin, 2015]), these methods have difficulty scaling up beyond approximately 5-6 dimensional state spaces.

An approach that is closely related to our work is the work presented in [Ny and Pappas, 2012]. The authors propose a randomized planning algorithm in the spirit of RRTs that explicitly reasons about disturbances and uncertainty. Specifications of input to output stability with respect to disturbances provide a parameterization of “tubes” (analogous to our “funnels”) that can be composed together to generate motion plans that are collision-free. The factors that distinguish the approach we present in this paper from the one proposed in [Ny and Pappas, 2012] are our focus on the *real-time* aspect of the problem and use of sums-of-squares programming as a way of computing reachable sets. In [Ny and Pappas, 2012], the focus is on generating safe motion plans when the obstacle positions are known *a priori*. Further, we provide a general technique for computing and explicitly minimizing the size of tubes.

Another approach that is closely related to ours is Model Predictive Control with Tubes [Mayne et al., 2005]. The idea is to solve the optimal control problem online with guaranteed “tubes” that the trajectories stay in. A closely related idea is that of “flow-tubes”, which have been used for high-level planning for autonomous systems [Li and Williams, 2008]. However, these methods are generally limited to linear systems.

2.4 Lyapunov Theory and SOS programming

A critical component of the work presented here is the computation of “funnels” for nonlinear systems via Lyapunov functions. The metaphor for thinking about Lyapunov functions as defining funnels was introduced to the robotics community in [Burridge et al., 1999], where funnels were *sequentially composed* in order to produce dynamic behaviors in a robot. However, computational tools for automatically searching for Lyapunov functions were lacking until very recently. In recent years, sums-of-squares (SOS) programming has emerged as a way of checking the Lyapunov function conditions associated with each funnel [Parrilo, 2000]. The technique relies on the ability to check nonnegativity of multivariate polynomials by expressing them as a sum of squares of polynomials. This can be written as a semidefinite optimization program and is amenable to efficient computational algorithms such as interior point methods [Parrilo, 2000]. Assuming polynomial dynamics, one can check that a polynomial Lyapunov candidate, $V(x)$, satisfies $V(x) > 0$ and $\dot{V}(x) < 0$ in some region B_r . Importantly, the same idea can be used for computing funnels along time-indexed trajectories of a system [Tedrake et al., 2010] [Tobenkin et al., 2011]. In this paper, we will use a similar approach to synthesize feedback controllers that explicitly seek to minimize the effect of disturbances on the system by minimizing the size of the funnel computed along a trajectory. Thus, we are guaranteed that if the system starts off in the set of given initial conditions, it will remain in the computed “funnel” even if the model of the dynamics is uncertain and the system is subjected to bounded disturbances.

The ability to compute funnels using SOS programming was leveraged by the LQR-Trees algorithm [Tedrake et al., 2010] for *feedback motion planning* for nonlinear systems. The algorithm works by creating a tree of locally stabilizing controllers which can take any initial condition in some bounded region in state space to the desired goal. However, LQR-Trees lack the ability to handle scenarios in which the task and environment are unknown till runtime: the offline precomputation of the tree does not take into account potential runtime constraints like obstacles, and an online implementation of the algorithm is computationally infeasible.

The SOS programming approach has also been used to guarantee obstacle avoidance conditions for nonlinear systems by computing *barrier certificates* [Prajna, 2006] [Barry et al., 2012]. Barrier functions are similar to Lyapunov functions in spirit, but are used to guarantee that trajectories starting in some given set of initial conditions will never enter an “unsafe” region containing obstacles. This approach, however, is limited to settings where the locations and geometry of obstacles are known beforehand since the barrier certificate one computes depends on this data and computing barrier certificates in real-time using SOS programming is not computationally feasible at the present time.

3 Background

In this section we provide a brief background on the key computational tools that will be employed throughout this paper.

3.1 Semidefinite Programming (SDP)

Semidefinite programs (SDPs) form an important class of convex optimization problems. They are optimization problems over the space of symmetric positive semidefinite (psd) matrices. Recall

that a $n \times n$ symmetric matrix Q is positive semidefinite if $x^T Q x \geq 0$, $\forall x \in \mathbb{R}^n$. Denoting the set of $n \times n$ symmetric matrices as \mathbf{S}^n , a SDP in standard form is written as:

$$\begin{aligned} \min_{X \in \mathbf{S}^n} \quad & \langle C, X \rangle \\ \text{s.t.} \quad & \langle A_i, X \rangle = b_i \quad \forall i \in \{1, \dots, m\}, \\ & X \succeq 0, \end{aligned} \tag{1}$$

where $C, A_i \in \mathbf{S}^n$ and $\langle X, Y \rangle := \text{Tr}(X^T Y) = \sum_{i,j} X_{ij} Y_{ij}$. In other words, a SDP involves minimizing a cost function that is linear in the elements of the decision matrix X subject to linear and positive semidefiniteness constraints on X .

Semidefinite programming includes Linear Programming (LP), Quadratic Programming (QP) and Second-Order Cone Programming (SOCP) as special cases. As in these other cases, SDPs are amenable to efficient numerical solution via interior point methods. The interested reader may wish to consult [Vandenberghe and Boyd, 1996] and [Blekherman et al., 2013, Section 2] for a more thorough introduction to SDPs.

3.2 Sums-of-Squares (SOS) Programming

An important application of SDPs is to check nonnegativity of polynomials. The decision problem associated with checking polynomial nonnegativity is NP-hard in general [Parrilo, 2000]. However, the problem of determining whether a polynomial is a sum-of-squares (SOS), which is a sufficient condition for nonnegativity, is amenable to efficient computation. A polynomial p in indeterminates¹ x_1, x_2, \dots, x_n is SOS if it can be written as $p(x) = \sum_{i=1}^m q_i^2(x)$ for a set of polynomials $\{q_i\}_{i=1}^m$. This condition is *equivalent* to the existence of a positive semidefinite (psd) matrix Q that satisfies:

$$p(x) = v(x)^T Q v(x), \quad \forall x \in \mathbb{R}^n, \tag{2}$$

where $v(x)$ is the vector of all monomials with degree less than or equal to half the degree of p [Parrilo, 2000]. Note that the equality constraint (2) imposes linear constraints on the elements of the matrix Q that come from matching coefficients of the polynomials on the left and right hand sides. Thus, semidefinite programming can be used to certify that a polynomial is a sum of squares. Indeed, by allowing the coefficients of the polynomial p to be decision variables, we can solve optimization problems over the space of SOS polynomials of some fixed degree. Such optimization problems are referred to as sums-of-squares (SOS) programs. The interested reader is referred to [Parrilo, 2000] and [Blekherman et al., 2013, Sections 3,4] for a more thorough introduction to SOS programming.

In addition to being able to prove global nonnegativity of polynomials, the SOS programming approach can also be used to demonstrate nonnegativity of polynomials on basic semialgebraic sets (i.e., sets described by a finite number of polynomial inequalities and equalities). Suppose we are given a set \mathcal{B} :

$$\mathcal{B} = \{x \in \mathbb{R}^n \mid g_{eq,i}(x) = 0, g_{ineq,j}(x) \geq 0\}, \tag{3}$$

¹Throughout this paper, the variables x that a polynomial $p(x)$ depend on will be referred to as “indeterminates”. This is to distinguish these variables from *decision variables* in our optimization problems, which will typically be the coefficients of the polynomial.

where $g_{eq,i}$ and $g_{ineq,j}$ are polynomials for $i \in \{1, \dots, N_{eq}\}, j \in \{1, \dots, N_{ineq}\}$. We are interested in showing that a polynomial p is nonnegative on the set \mathcal{B} :

$$x \in \mathcal{B} \implies p(x) \geq 0. \quad (4)$$

We can write the following SOS constraints in order to impose (4):

$$q(x) := p(x) - \underbrace{\sum_{i=1}^{N_{eq}} L_{eq,i}(x)g_{eq,i}(x)}_{r(x)} - \sum_{j=1}^{N_{ineq}} L_{ineq,j}(x)g_{ineq,j}(x) \text{ is SOS,} \quad (5)$$

$$L_{ineq,j}(x) \text{ is SOS, } \forall j \in \{0, \dots, N_{ineq}\}. \quad (6)$$

Here, the polynomials $L_{eq,i}$ and $L_{ineq,j}$ are “multiplier” polynomials analogous to Lagrange multipliers in constrained optimization. In order to see that (5) and (6) are sufficient conditions for (4), note that when a point x satisfies $g_{eq,i}(x) = 0$ and $g_{ineq,j}(x) \geq 0$ for $i \in \{1, \dots, N_{eq}\}, j \in \{1, \dots, N_{ineq}\}$ (i.e., when $x \in \mathcal{B}$) then the term $r(x)$ is non-positive. Hence, for $q(x)$ to be nonnegative (which must be the case since q is SOS), $p(x)$ must be nonnegative. Thus, we have the desired implication in (4). This process for using multipliers to impose nonnegativity constraints on sets is known as the generalized S-procedure [Parrilo, 2000] and will be used extensively in Section 4 for computing funnels.

4 Computing Funnels

In this section we describe how the tools from Section 3 can be used to compute outer approximations of reachable sets (“funnels”) around trajectories of a nonlinear system. The approach in Section 4.1 builds on the work presented in [Tobenkin et al., 2011, Tedrake et al., 2010] while Sections 4.3.1 and 4.3.2 are based on [Majumdar and Tedrake, 2012] and [Majumdar et al., 2013] respectively. In contrast to this prior work however, we consider the problem of computing outer approximations of forwards reachable sets as opposed to inner approximations of backwards reachable sets. This leads to a few subtle differences in the cost functions of our optimization problems.

Consider the following dynamical system:

$$\dot{x} = f(x(t), u(t)), \quad (7)$$

where $x(t) \in \mathbb{R}^n$ is the state of the system at time t and $u(t) \in \mathbb{R}^m$ is the control input. Let $x_0 : [0, T] \rightarrow \mathbb{R}^n$ be the nominal trajectory that we would like the system to follow and $u_0 : [0, T] \rightarrow \mathbb{R}^m$ be the nominal open-loop control input. Defining new coordinates $\bar{x} = x - x_0(t)$ and $\bar{u} = u - u_0(t)$, we can rewrite the dynamics (7) in these variables as:

$$\dot{\bar{x}} = \dot{x} - \dot{x}_0 = f(x_0(t) + \bar{x}(t), u_0(t) + \bar{u}(t)) - \dot{x}_0. \quad (8)$$

We will first consider the problem of computing funnels for a closed-loop system subject to no uncertainty. To this end, we assume that we are given a feedback controller $\bar{u}_f(t, \bar{x})$ that corrects for deviations around the nominal trajectory (we will consider the problem of designing feedback controllers later in this section). We can then write the closed-loop dynamics of the system as:

$$\dot{\bar{x}} = f_{cl}(t, \bar{x}(t)). \quad (9)$$

Given a set of initial conditions $\mathcal{X}_0 \subset \mathbb{R}^n$ with $x_0(0) \in \mathcal{X}_0$, our goal is to find a tight outer approximation of the set of states the system may evolve to at time $t \in [0, T]$. In particular, we are concerned with finding sets $F(t) \subset \mathbb{R}^n$ such that:

$$\bar{x}(0) \in \mathcal{X}_0 \implies \bar{x}(t) \in F(t), \quad \forall t \in [0, T]. \quad (10)$$

Definition 1. A funnel associated with a closed-loop dynamical system $\dot{x} = f_{cl}(t, \bar{x}(t))$ is a map $F : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^n)$ from the time-interval $[0, T]$ to the power set (i.e., the set of subsets) of \mathbb{R}^n such that the sets $F(t)$ satisfy the condition (10) above.

Thus, with each time $t \in [0, T]$, the funnel associates a set $F(t) \subset \mathbb{R}^n$. We will parameterize the sets $F(t)$ as sub-level sets of nonnegative time-varying functions $V : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^+$:

$$F(t) = \{\bar{x}(t) \in \mathbb{R}^n \mid V(t, \bar{x}(t)) \leq \rho(t)\}. \quad (11)$$

Letting $\mathcal{X}_0 \subset F(0, \bar{x})$, the following constraint is a sufficient condition for (10):

$$V(t, \bar{x}) = \rho(t) \implies \dot{V}(t, \bar{x}) < \dot{\rho}(t), \quad \forall t \in [0, T]. \quad (12)$$

Here, \dot{V} is computed as:

$$\dot{V}(t, \bar{x}) = \frac{\partial V(t, \bar{x})}{\partial \bar{x}} f_{cl}(t, \bar{x}) + \frac{\partial V(t, \bar{x})}{\partial t}. \quad (13)$$

Intuitively, the constraint (12) says that on the boundary of the funnel (i.e., when $V(t, \bar{x}) = \rho(t)$), the function V grows slower than ρ . Hence, states on the boundary of the funnel remain within the funnel. This intuition is formalized in [Tedrake et al., 2010, Tobenkin et al., 2011].

While *any* function that satisfies (12) provides us with a valid funnel, we are interested in finding *tight* outer approximations of the reachable set. A natural cost function for encouraging tightness is the volume of the sets $F(t)$. Combining this cost function with our constraints, we obtain the following optimization problem:

$$\begin{aligned} \inf_{V, \rho} \quad & \int_0^T \text{vol}(F(t)) \, dt \\ \text{s.t.} \quad & V(t, \bar{x}) = \rho(t) \implies \dot{V}(t, \bar{x}) < \dot{\rho}(t), \quad \forall t \in [0, T], \\ & \mathcal{X}_0 \subset F(0, \bar{x}). \end{aligned} \quad (14)$$

4.1 Numerical implementation using SOS programming

Since the optimization problem (14) involves searching over spaces of functions, it is infinite dimensional and hence not directly amenable to numerical computation. However, we can use the SOS programming approach described in Section 3 to obtain finite dimensional optimization problems in the form of semidefinite programs (SDPs). We first concentrate on implementing the constraints in (14). We will assume that the initial condition set \mathcal{X}_0 is a semi-algebraic set (i.e., described in terms of polynomial inequalities):

$$\mathcal{X}_0 = \{\bar{x} \in \mathbb{R}^n \mid g_{0,i}(\bar{x}) \geq 0, \quad \forall i = 1, \dots, N_0\}. \quad (15)$$

Then the constraints in (14) can be written as:

$$V(t, \bar{x}) = \rho(t) \implies \rho(t) - \dot{V}(t, \bar{x}) > 0 \quad (16)$$

$$g_{0,i}(\bar{x}) \geq 0 \ \forall i \in \{1, \dots, N_0\} \implies \rho(0) - V(0, \bar{x}) \geq 0. \quad (17)$$

If we restrict ourselves to polynomial dynamics and polynomial functions V and ρ , these constraints are precisely in the form of (4) in Section 3.2. We can thus apply the procedure described in Section 3.2 and arrive at the following sufficient conditions for (16) and (17):

$$\dot{\rho}(t) - \dot{V}(t, \bar{x}) - L(t, \bar{x})[V(t, \bar{x}) - \rho(t)] - L_t(t, \bar{x})[t(T - t)] \text{ is SOS,} \quad (18)$$

$$\rho(0) - V(0, \bar{x}) - \sum_i^{N_0} L_{0,i}(\bar{x})g_{0,i}(\bar{x}) \text{ is SOS,} \quad (19)$$

$$L_t(t, \bar{x}), L_{0,i}(\bar{x}) \text{ are SOS, } \forall i \in \{1, \dots, N_0\}. \quad (20)$$

As in Section 3.2, the polynomials L, L_t and $L_{0,i}$ are “multiplier” polynomials whose coefficients are decision variables.

Next, we focus on approximating the cost function in (14) using semidefinite programming. This can be achieved by sampling in time and replacing the integral with the finite sum $\sum_{k=1}^N \text{vol}(F(t_k))$. In the special case where the function V is quadratic in \bar{x} :

$$V(t_k, \bar{x}) = \bar{x}^T S_k \bar{x}, \quad S_k \succeq 0, \quad (21)$$

the set $F(t_k)$ is an ellipsoid and we can use semidefinite programming (SDP) to directly minimize the volume by maximizing the determinant of S_k (recall that the volume of the ellipsoid $F(t_k)$ is a monotonically decreasing function of the determinant of S_k). Note that while the problem of maximizing the determinant of a psd matrix is not directly a problem of the form (1), it can be transformed into such a form [Ben-Tal and Nemirovski, 2001, Section 3]. Further note that the fact that our cost function can be handled directly in the SDP framework is in distinction to the approaches for computing inner approximations of backwards reachable sets [Tedrake et al., 2010] [Tobenkin et al., 2011] [Majumdar et al., 2013]. This is because the determinant of a psd matrix is a *concave* function and hence *minimizing* the determinant is not a convex problem. Hence, in the previous work, the authors used heuristics for maximizing volume.

In the more general case, we can minimize an upper bound on the cost function $\sum_{k=1}^N \text{vol}(F(t_k))$ by introducing ellipsoids $\mathcal{E}(t_k)$:

$$\mathcal{E}(t_k) = \{\bar{x} \in \mathbb{R}^n \mid \bar{x}^T S_k \bar{x} \leq 1, \quad S_k \succeq 0\} \quad (22)$$

such that $F(t_k) \subset \mathcal{E}(t_k)$ and minimizing $\sum_{k=1}^N \text{vol}(\mathcal{E}(t_k))$. The containment constraint can be equivalently expressed as the constraint:

$$V(t_k, \bar{x}) \leq \rho(t_k) \implies \bar{x}^T S_k \bar{x} \leq 1, \quad (23)$$

and can thus be imposed using SOS constraints:

$$1 - \bar{x}^T S_k \bar{x} - L_{\mathcal{E},k}(\bar{x})[\rho(t_k) - V(t_k, \bar{x})] \text{ is SOS,} \quad (24)$$

$$L_{\mathcal{E},k}(\bar{x}) \text{ is SOS.} \quad (25)$$

Combining our cost function with the constraints (18) - (20), we obtain the following optimization problem:

$$\inf_{V, \rho, L, L_t, L_{0,i}, S_k, L_{\mathcal{E},k}} \sum_{k=1}^N \text{vol}(\mathcal{E}(t_k)) = \sum_{k=1}^N \text{vol}(\{\bar{x} | \bar{x}^T S_k \bar{x} \leq 1\}) \quad (26)$$

$$\text{s.t.} \quad \dot{\rho}(t) - \dot{V}(t, \bar{x}) - L(t, \bar{x})[V(t, \bar{x}) - \rho(t)] - L_t(t, \bar{x})[t(T - t)] \text{ is SOS,} \quad (27)$$

$$\rho(0) - V(0, \bar{x}) - \sum_i^{N_0} L_{0,i}(\bar{x})g_{0,i}(\bar{x}) \text{ is SOS,}$$

$$1 - \bar{x}^T S_k \bar{x} - L_{\mathcal{E},k}(\bar{x})[\rho(t_k) - V(t_k, \bar{x})] \text{ is SOS,} \quad \forall k \in \{1, \dots, N\},$$

$$S_k \succeq 0, \quad \forall k \in \{1, \dots, N\},$$

$$L_t(t, \bar{x}), L_{0,i}(\bar{x}), L_{\mathcal{E},k}(\bar{x}) \text{ are SOS,} \quad \forall i \in \{1, \dots, N_0\}, \quad \forall k \in \{1, \dots, N\}.$$

While this optimization problem is finite dimensional, it is non-convex in general since the first constraints are *bilinear* in the decision variables (e.g., the coefficients of the polynomials L and V are multiplied together in the first constraint). To apply SOS programming, we require the constraints to be linear in the coefficients of the polynomials we are optimizing. However, note that when V and ρ are fixed, the constraints are linear in the other decision variables. Similarly, when the multipliers L and $L_{\mathcal{E},k}$ are fixed, the constraints are linear in the remaining decision variables. Thus, we can efficiently perform this optimization by alternating between the two sets of decision variables $(L, L_t, L_{0,i}, S_k, L_{\mathcal{E},k})$ and $(V, \rho, L_t, L_{0,i}, S_k)$. In each step of the alternation, we can optimize our cost function $\sum_{k=1}^N \text{vol}(\mathcal{E}(t_k))$. These alternations are summarized in Algorithm 1. Note that the algorithm requires an initialization for V and ρ . We will discuss how to obtain these in Section 4.4.

- 1: Initialize V and ρ .
- 2: $\text{cost}_{\text{prev}} = \infty$;
- 3: $\text{converged} = \text{false}$;
- 4: **while** $\neg \text{converged}$ **do**
- 5: **STEP 1** : Minimize $\sum_{k=1}^N \text{vol}(\mathcal{E}(t_k))$ by searching for multiplier polynomials $(L, L_t, L_{0,i}, L_{\mathcal{E},k})$ and S_k while fixing V and ρ .
- 6: **STEP 2** : Minimize $\sum_{k=1}^N \text{vol}(\mathcal{E}(t_k))$ by searching for $(V, \rho, L_t, L_{0,i}, S_k)$ while fixing L and $L_{\mathcal{E},k}$.
- 7: $\text{cost} = \sum_{k=1}^N \text{vol}(\mathcal{E}(t_k))$;
- 8: **if** $\frac{\text{cost}_{\text{prev}} - \text{cost}}{\text{cost}_{\text{prev}}} < \epsilon$ **then**
- 9: $\text{converged} = \text{true}$;
- 10: **end if**
- 11: $\text{cost}_{\text{prev}} = \text{cost}$;
- 12: **end while**

Algorithm 1: Funnel Computation

Remark 1. It is easy to see that Algorithm 1 converges (though not necessarily to an optimal solution). Each iteration of the alternations is guaranteed to achieve a cost function that is at least as good as the previous iteration (since the solution from the previous iteration is a valid one). Hence, the sequence of optimal values in each iteration form a monotonically non-increasing

sequence. Combined with the fact that the cost function is bounded below by 0, we conclude that this sequence converges and hence that Algorithm 1 terminates.

4.2 Approximation via time-sampling

As observed in [Tobenkin et al., 2011] in practice it is often the case that the nominal trajectory $x_0 : [0, T] \rightarrow \mathbb{R}^n$ is difficult to approximate with a low degree polynomial in time. This can lead to the constraint (27) in the problem (26) having a high degree polynomial dependence on t . Thus it is often useful to implement an approximation of the optimization problem (26) where the condition (16) is checked only at a finite number of sample points $t_k \in [0, T]$, $k \in \{1, \dots, N\}$. We can use a piecewise linear parameterization of ρ and can thus compute:

$$\dot{\rho}(t_k) = \frac{\rho(t_{k+1}) - \rho(t_k)}{t_{k+1} - t_k}. \quad (28)$$

Similarly we can parameterize the function V by polynomials $V_k(\bar{x})$ at each time sample and compute:

$$\frac{\partial V(t, \bar{x})}{\partial t} \approx \frac{V_{k+1}(\bar{x}) - V_k(\bar{x})}{t_{k+1} - t_k}. \quad (29)$$

We can then write the following modified version of the problem (26):

$$\begin{aligned} \inf_{V_k, \rho, L_k, L_{0,i}, S_k, L_{\mathcal{E},k}} \quad & \sum_{k=1}^N \text{vol}(\mathcal{E}(t_k)) = \sum_{k=1}^N \text{vol}(\{\bar{x} | \bar{x}^T S_k \bar{x} \leq 1\}) \\ \text{s.t.} \quad & \dot{\rho}(t_k) - \dot{V}_k(\bar{x}) - L_k(\bar{x})[V_k(\bar{x}) - \rho(t_k)], \quad \forall k \in \{1, \dots, N\}, \\ & \rho(t_1) - V_1(\bar{x}) - \sum_i^{N_0} L_{0,i}(\bar{x}) g_{0,i}(\bar{x}) \text{ is SOS}, \\ & 1 - \bar{x}^T S_k \bar{x} - L_{\mathcal{E},k}(\bar{x})[\rho(t_k) - V_k(\bar{x})] \text{ is SOS}, \quad \forall k \in \{1, \dots, N\}, \\ & S_k \succeq 0, \quad \forall k \in \{1, \dots, N\}, \\ & L_{0,i}(\bar{x}), L_{\mathcal{E},k}(\bar{x}) \text{ are SOS}, \quad \forall i \in \{1, \dots, N_0\}, \forall k \in \{1, \dots, N\}. \end{aligned} \quad (30)$$

This program does not have any algebraic dependence on the variable t and can thus provide significant computational gains over (26). However, it does not provide an exact funnel certificate. One would hope that with a sufficiently fine sampling in time, one would recover exactness. Partial results in this direction are provided in [Tobenkin et al., 2011] along with numerical examples showing that the loss of accuracy from the sampling approximation can be quite small in practice.

The problem (30) is again bilinear in the decision variables. It is linear in the two sets of decision variables $(L_k, L_{0,i}, S_k, L_{\mathcal{E},k})$ and $(V_k, \rho, L_{0,i}, S_k)$. Thus, Algorithm 1 can be applied directly to (30) with the minor modification that V and ρ are replaced by their time-sampled counterparts and the multipliers (L, L_t) are replaced by the multipliers L_k .

4.3 Extensions to the basic algorithm

Next we describe several extensions to the basic framework for computing funnels described in Section 4.1. Section 4.3.1 discusses the scenario in which the dynamics of the system are subject to bounded disturbances/uncertainty, Section 4.3.2 considers the problem of synthesizing feedback

controllers that explicitly attempt to minimize the size of the funnel, Section 4.3.3 demonstrates how to handle input saturations, and Section 4.3.4 considers a generalization of the cost function.

4.3.1 Uncertainty in the dynamics

Suppose that the dynamics of the system are subject to an uncertainty term $w(t) \in \mathbb{R}^d$ that models external disturbances or parametric model uncertainties. The closed-loop dynamics (9) can then be modified to capture this uncertainty:

$$\dot{\bar{x}} = f_{cl}(t, \bar{x}(t), w(t)). \quad (31)$$

We will assume that the dynamics f_{cl} depend polynomially on w . Given an initial condition set $\mathcal{X}_0 \subset \mathbb{R}^n$ as before, our goal is to find sets $F(t)$ such that $\bar{x}(t)$ is guaranteed to be in $F(t)$ for any valid disturbance profile:

$$\bar{x}(0) \in \mathcal{X}_0 \implies \bar{x}(t) \in F(t), \forall t \in [0, T], \forall w : [0, T] \rightarrow \mathcal{W}. \quad (32)$$

Parameterizing the sets $F(t)$ as sub-level sets of nonnegative time-varying functions $V : [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^+$ as before, the following condition is sufficient to ensure (32):

$$V(t, \bar{x}) = \rho(t) \implies \dot{V}(t, \bar{x}, w) < \dot{\rho}(t), \forall t \in [0, T], \forall w(t) \in \mathcal{W}, \quad (33)$$

where \dot{V} is computed as:

$$\dot{V}(t, \bar{x}, w) = \frac{\partial V(t, \bar{x})}{\partial \bar{x}} f_{cl}(t, \bar{x}, w) + \frac{\partial V(t, \bar{x})}{\partial t}. \quad (34)$$

This is almost identical to the condition (12), with the exception that the function V is required to decrease on the boundary of the funnel for every choice of disturbance. Assuming that the set \mathcal{W} is a semi-algebraic set $\mathcal{W} = \{w \in \mathbb{R}^d \mid g_{w,j}(w) \geq 0, \forall j = 1, \dots, N_w\}$, the optimization problem (26) is then easily modified by replacing condition (27) with the following constraints:

$$\begin{aligned} \dot{\rho}(t) - \dot{V}(t, \bar{x}, w) - L(t, \bar{x}, w)[V(t, \bar{x}) - \rho(t)] - L_t(t, \bar{x}, w)[t(T - t)] \dots \\ \dots - \sum_{j=1}^{N_w} L_{w,j}(t, \bar{x}, w)g_{w,j}(w) \text{ is SOS,} \\ L_{w,j}(t, \bar{x}, w) \text{ is SOS, } \forall j = \{1, \dots, N_w\}. \end{aligned} \quad (35)$$

These SOS constraints now involve polynomials in the indeterminates t, \bar{x} and w . Since these constraints are linear in the coefficients of the newly introduced multipliers $L_{w,j}$, Algorithm 1 can be directly applied to the modified optimization problem by adding $L_{w,j}$ to the list of polynomials to be searched for in both Step 1 and Step 2 of the iterations. Similarly, the time-sampled approximation described in Section 4.2 can also be applied to (35).

4.3.2 Feedback control synthesis

So far we have assumed that we have been provided with a feedback controller that corrects for deviations around the nominal trajectory. We now consider the problem of *optimizing* the feedback controller in order to minimize the size of the funnel. We will assume that the system is control affine:

$$\dot{x} = f(x(t)) + g(x(t))u(t), \quad (36)$$

and parameterize the control policy as a polynomial $\bar{u}_f(t, \bar{x})$. We can thus write the dynamics in the \bar{x} coordinates as:

$$\dot{\bar{x}} = f(x_0(t) + \bar{x}(t)) + g(x(t))[u_0(t) + \bar{u}_f(t, \bar{x})] - \dot{x}_0. \quad (37)$$

The feedback controller can then be optimized by adding the coefficients of the polynomial $\bar{u}_f(t, \bar{x})$ to the set of decision variables in the optimization problem (26) while keeping all the constraints unchanged. Note that \bar{u}_f appears in the constraints only through \dot{V} , which is now bilinear in the coefficients of V and \bar{u}_f since:

$$\dot{V}(t, \bar{x}) = \frac{\partial V(t, \bar{x})}{\partial \bar{x}} \dot{\bar{x}} + \frac{\partial V(t, \bar{x})}{\partial t}. \quad (38)$$

With the (coefficients of) the feedback controller \bar{u}_f as part of the optimization problem, note that the constraints of the problem (26) are now bilinear in the two sets of decision variables $(L, L_t, L_{0,i}, S_k, L_{\mathcal{E},k}, \bar{u}_f)$ and $(V, \rho, L_t, L_{0,i}, S_k)$. Thus, in principle we could use a bilinear alternation scheme similar to the one in Algorithm 1 and alternatively optimize these two sets of decision variables. However, in this case we would not be searching for a controller that explicitly seeks to minimize the size of the funnel (since the controller would not be searched for at the same time as V or ρ , which define the funnel). To get around this issue, we add another step in each iteration where we optimize our cost function $\sum_{k=1}^N \text{vol}(\mathcal{E}(t_k))$ by searching for $(\bar{u}_f, \rho, L_t, L_{0,i}, S_k)$ while keeping $(V, L, L_{\mathcal{E},k})$ fixed. This allows us to search for \bar{u}_f and ρ at the same time, which can significantly improve the quality of the controllers and funnels we obtain. These steps are summarized in Algorithm 2. By a reasoning identical to the one in Remark 1 it is easy to see that the sequence of optimal values produced by Algorithm 2 converges.

We note that the approach for taking into account uncertainty described in Section 4.3.1 can easily be incorporated into Algorithm 2. Similarly, by parameterizing the controller \bar{u}_f as polynomials $u_{f,k}(\bar{x})$ at the times t_k , we can also apply the time-sampled approximation described in Section 4.2.

4.3.3 Actuator saturations

A. Approach 1

Our approach also allows us to incorporate actuator limits into the verification procedure. Although we examine the single-input case in this section, this framework is easily extended to handle multiple inputs. Let the control input $u(t)$ at time t be mapped through the following control saturation function:

$$s(u(t)) = \begin{cases} u_{\max} & \text{if } u(t) \geq u_{\max} \\ u_{\min} & \text{if } u(t) \leq u_{\min} \\ u(t) & \text{o.w.} \end{cases}$$

```

1: Initialize  $V$  and  $\rho$ .
2:  $cost_{prev} = \infty$ ;
3: converged = false;
4: while  $\neg$ converged do
5:   STEP 1 : Minimize  $\sum_{k=1}^N vol(\mathcal{E}(t_k))$  by searching for controller  $\bar{u}_f$  and  $(L, L_t, L_{0,i}, L_{\mathcal{E},k}, S_k)$  while fixing  $V$  and  $\rho$ .
6:   STEP 2 : Minimize  $\sum_{k=1}^N vol(\mathcal{E}(t_k))$  by searching for controller  $\bar{u}_f$  and  $(\rho, L_t, L_{0,i}, S_k)$  while fixing  $(V, L, L_{\mathcal{E},k})$ .
7:   STEP 3 : Minimize  $\sum_{k=1}^N vol(\mathcal{E}(t_k))$  by searching for  $(V, \rho, L_t, L_{0,i}, S_k)$  while fixing  $(L, L_{\mathcal{E},k}, \bar{u}_f)$ .
8:    $cost = \sum_{k=1}^N vol(\mathcal{E}(t_k))$ ;
9:   if  $\frac{cost_{prev} - cost}{cost_{prev}} < \epsilon$  then
10:    converged = true;
11:   end if
12:    $cost_{prev} = cost$ ;
13: end while

```

Algorithm 2: Feedback Control Synthesis

Then, in a manner similar to [Tedrake et al., 2010], a piecewise analysis of $\dot{V}(t, \bar{x})$ can be used to check the Lyapunov conditions are satisfied even when the control input saturates. Defining:

$$\dot{V}_{min}(t, \bar{x}) := \frac{\partial V(t, \bar{x})}{\partial \bar{x}} \left(f(x_0 + \bar{x}) + g(x_0 + \bar{x})u_{min} \right) + \frac{\partial V(t, \bar{x})}{\partial t}, \quad (39)$$

$$\dot{V}_{max}(t, \bar{x}) := \frac{\partial V(t, \bar{x})}{\partial \bar{x}} \left(f(x_0 + \bar{x}) + g(x_0 + \bar{x})u_{max} \right) + \frac{\partial V(t, \bar{x})}{\partial t}, \quad (40)$$

we must check the following conditions:

$$u_0(t) + \bar{u}_f(t, \bar{x}) \leq u_{min}, \quad V(t, \bar{x}) = \rho(t) \implies \dot{V}_{min}(t, \bar{x}) < \dot{\rho}(t), \quad (41)$$

$$u_0(t) + \bar{u}_f(t, \bar{x}) \geq u_{max}, \quad V(t, \bar{x}) = \rho(t) \implies \dot{V}_{min}(t, \bar{x}) < \dot{\rho}(t), \quad (42)$$

$$u_{min} \leq u_0(t) + \bar{u}_f(t, \bar{x}) \leq u_{max}, \quad V(t, \bar{x}) = \rho(t) \implies \dot{V}(t, \bar{x}) < \dot{\rho}(t), \quad (43)$$

where u_0 is the open-loop control input and \bar{u}_f is the feedback controller as before. These conditions can be enforced by adding additional multipliers to the optimization program (26) or its time-sampled counterpart (30).

B. Approach 2

Although one can handle multiple inputs via the above method, the number of SOS conditions grows exponentially with the number of inputs (3^m conditions for \dot{V} are needed in general to handle all possible combinations of input saturations). Thus, for systems with a large number of inputs, an alternative formulation was proposed in [Majumdar et al., 2013] that avoids this exponential growth in the size of the SOS program at the cost of adding conservativeness to the size of the funnel. Given limits on the control vector $u \in \mathbb{R}^m$ of the form $u_{min} < u < u_{max}$, we can ask to satisfy:

$$\bar{x} \in F(t) \implies u_{min} < u_0(t) + \bar{u}_f(t, \bar{x}) < u_{max}, \quad \forall t \in [0, T]. \quad (44)$$

This constraint implies that the applied control input remains within the specified bounds inside the verified funnel (a conservative condition). The number of extra constraints grows linearly with the number of inputs (since we have one new condition for every input), thus leading to smaller optimization problems.

4.3.4 A more general cost function

We end our discussion of extensions to the basic algorithm for computing funnels presented in Section 4.1 by considering a generalization of the cost function (volume of the funnel) we have used so far. In particular, it is sometimes useful to minimize the volume of the funnel *projected* onto a subspace of the state space. Suppose this projection map is given by $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^{n_p}$ with a corresponding $n_p \times n$ projection matrix P . For an ellipsoid $\mathcal{E} = \{\bar{x} \in \mathbb{R}^n \mid \bar{x}^T S_k \bar{x} \leq 1\}$, the projected set $\pi(\mathcal{E})$ is also an ellipsoid $\mathcal{E}_p = \{\bar{x} \in \mathbb{R}^{n_p} \mid \bar{x}^T S_k^{(p)} \bar{x} \leq 1\}$ with:

$$S_k^{(p)} = (PS_k^{-1}P^T)^{-1}. \quad (45)$$

Recall that the ability to minimize the volume of the ellipsoid \mathcal{E} using SDP relied on being able to maximize the determinant of S_k . In order to minimize the volume of \mathcal{E}_p , we would have to maximize $\det(S_k^{(p)})$, which is a complicated (i.e. nonlinear) function of S_k . Hence, in each iteration of Algorithm 1 we linearize the function $\det(S_k^{(p)})$ with respect to S_k at the solution of S_k from the previous iteration and maximize this linearization instead. The linearization of $\det(S_k^{(p)})$ with respect to S_k at a nominal value $S_{k,0}$ can be explicitly computed as:

$$\text{Tr}\left(P^T(PS_{k,0}^{-1}P^T)^{-1}PS_{k,0}^{-1}S_kS_{k,0}^{-1}\right), \quad (46)$$

where Tr refers to the trace of the matrix.

4.4 Implementation details

We end this section on computing funnels by discussing a few important implementation details.

4.4.1 Trajectory generation

An important step that is necessary for the success of our approach to computing funnels is the generation of a dynamically feasible open-loop control input $u_0 : [0, T] \mapsto \mathbb{R}^m$ and corresponding nominal trajectory $x_0 : [0, T] \mapsto \mathbb{R}^n$. A method that has been shown to work well in practice and scale to high dimensions is the direct collocation trajectory optimization method [Betts, 2001]. While this is the approach we use for the results in Section 7, other methods like the Rapidly Exploring Randomized Tree (RRT) or its asymptotically optimal version, RRT* can be used too [Kuffner and Lavalle, 2000, Karaman and Frazzoli, 2011].

4.4.2 Initializing V and ρ

Algorithms 1 and 2 require an initial guess for the functions V and ρ . In [Tedrake et al., 2010], the authors use the Lyapunov function candidate associated with a time-varying LQR controller. The control law is obtained by solving a Riccati differential equation:

$$-\dot{S}(t) = Q - S(t)B(t)R^{-1}B^T S(t) + S(t)A(t) + A(t)^T S(t) \quad (47)$$




Figure 2: The ordered pair of funnels (F_1, F_2) is sequentially composable. The outlet of F_1 is contained within the inlet of F_2 , i.e., $F_1(T_1) \subset F_2(0)$.

with final value conditions $S(t) = S_f$. Here $A(t)$ and $B(t)$ describe the time-varying linearization of the dynamics about the nominal trajectory x_0 . The matrices Q and R are positive-definite cost-matrices. The function:

$$V_{\text{guess}}(t, \bar{x}) = (x - x_0(t))^T S(t)(x - x_0(t)) = \bar{x}^T S(t) \bar{x} \quad (48)$$

is our initial Lyapunov candidate. Setting ρ to a quickly increasing function such as an exponential is typically sufficient to obtain a feasible initialization.

5 Funnel Libraries

5.1 Sequential composition

One can think of funnels computed using the machinery described in Section 4 as *robust* motion primitives (the robustness is to initial conditions and uncertainty in the dynamics). While we could define a *funnel library* simply as a collection \mathcal{F} of funnels and associated feedback controllers, it will be fruitful to associate some additional structure with \mathcal{F} . In particular, it is useful to know how funnels can be *sequenced* together to form composite robust motion plans. In order to consider this more formally, we will first introduce the notion of *sequential composition* of funnels defined in [Burridge et al., 1999].

Definition 2. [Burridge et al., 1999] *An ordered pair (F_1, F_2) of funnels $F_1 : [0, T_1] \rightarrow \mathcal{P}(\mathbb{R}^n)$ and $F_2 : [0, T_2] \rightarrow \mathcal{P}(\mathbb{R}^n)$ is sequentially composable if $F_1(T_1) \subset F_2(0)$.*

In other words, two funnels are sequentially composable if the “outlet” of one is contained within the “inlet” of the other. A pictorial depiction of this is provided in Figure 2. We note that the sequential composition of two such funnels is itself a funnel.

5.2 Exploiting invariances in the dynamics

For our purposes here, it is useful to introduce a slightly generalized notion of sequential compossibility that will allow us to exploit *invariances* (continuous symmetries) in the dynamics. In

particular, the dynamics of large classes of mechanical systems such as mobile robots are often invariant under certain transformations of the state space. For Lagrangian systems, the notion of “cyclic coordinates” captures such invariances. A cyclic coordinate is a (generalized) coordinate of the system that the Lagrangian does *not* depend on. We can then write the dynamics of the system $\dot{x} = f(x(t), u(t))$ as a function of a state vector $x = [x_c, x_{nc}]$ which is partitioned into cyclic coordinates x_c and non-cyclic coordinates x_{nc} in such a way that the dynamics only depend on the non-cyclic coordinates:

$$\dot{x} = f(x_{nc}(t), u(t)). \quad (49)$$

For example, the dynamics of a quadrotor or fixed-wing airplane (expressed in an appropriate coordinate system) do not depend on the $x - y - z$ position of the system or the yaw angle.

Invariance of the dynamics of the system also implies that if a curve $t \mapsto (x(t), u(t))$ is a valid solution to the dynamics $\dot{x} = f(x(t), u(t))$, then so is the transformed solution:

$$t \mapsto (\Psi_c(x(t)), u(t)), \quad (50)$$

where Ψ_c is a transformation of the state vector along the cyclic coordinates. This allows us to make the following important observation.

Remark 2. *Suppose we are given a system whose dynamics are invariant to transformations Ψ_c along cyclic coordinates x_c . Let $F : [0, T] \rightarrow \mathcal{P}(\mathbb{R}^n)$ given by $t \mapsto F(t)$ be a funnel associated with this system. Then the transformed funnel given by $t \mapsto \Psi_c(F(t))$ is also a valid funnel. Hence, one can in fact think of invariances in the dynamics giving rise to an infinite family of funnels parameterized by shifts $\Psi_c(F)$ of a funnel F along cyclic coordinates of the system.*

Note that here we have implicitly assumed that the feedback controller:

$$u_f(t, x) = u_0(t) + \bar{u}_f(t, \bar{x}) = u_0(t) + \bar{u}_f(t, x - x_0(t)) \quad (51)$$

associated with the funnel has also been transformed to:

$$u_0(t) + \bar{u}_f(t, x - \Psi_c(x_0(t))). \quad (52)$$

In other words, we have transformed the reference trajectory we are tracking by Ψ_c . Henceforth, when we refer to transformations of funnels along cyclic coordinates we will implicitly assume that the feedback controller has also been appropriately modified in this manner.

These observations allow us to define a generalized notion of sequential composition that exploits invariances in the dynamics. We will refer to this notion as sequential composition *modulo invariances (MI)*.

Definition 3. *An ordered pair (F_1, F_2) of funnels $F_1 : [0, T_1] \rightarrow \mathcal{P}(\mathbb{R}^n)$ and $F_2 : [0, T_2] \rightarrow \mathcal{P}(\mathbb{R}^n)$ is sequentially composable modulo invariances (MI) if there exists a transformation Ψ_c of the state along cyclic coordinates such that $F_1(T_1) \subset \Psi_c(F_2(0))$.*

Informally, two funnels F_1 and F_2 are sequentially composable in this generalized sense if one can shift F_2 along the cyclic coordinates of the system and ensure that its inlet contains the outlet of F_1 . Figure 3 provides a pictorial depiction of this.




Figure 3: Sequential composition modulo invariances. The top row of the figure shows two funnels that are *not* sequentially composable in the sense of Definition 2. However, as shown in the bottom row of the figure, they *are* sequentially composable in the more general sense of Definition 3. By shifting the funnel F_2 (whose outline is plotted using dotted lines for reference) in the cyclic coordinate (x_c), we can ensure that the outlet of F_1 lies in the inlet of this shifted funnel $\Psi_c(F_2)$.

One may think of sequential composability MI of funnels as being analogous to the compatibility condition required for sequencing trajectories in the library of a Maneuver Automaton [Frazzoli et al., 2005]. Let π_{nc} denote the projection operator that maps a state $x = [x_c, x_{nc}]$ to the non-cyclic coordinates x_{nc} . In order to be able to sequence together two trajectories $x_1 : [0, T_1] \rightarrow \mathbb{R}^n$ and $x_2 : [0, T_2] \rightarrow \mathbb{R}^n$, one requires:

$$\pi_{nc}(x_1(T_1)) = \pi_{nc}(x_2(0)). \quad (53)$$

Note, however that imposing this compatibility condition on the nominal trajectories associated with two funnels is neither necessary nor sufficient for the funnels being sequentially composable MI. Sequentially composable MI is concerned with the compatibility between funnels themselves and not the underlying nominal trajectories, which distinguishes our notion of compatibility between maneuvers from that of [Frazzoli et al., 2005].

5.3 Runtime composability

The two notions of sequential composability we have considered so far allow us to produce new funnels from a given set of funnels by stitching them together appropriately in an offline pre-processing stage. We now introduce another notion of composability that is particularly important for reasoning about how funnels can be executed sequentially at runtime. We will refer to this notion as *runtime composability*.

Definition 4. *An ordered pair (F_1, F_2) of funnels $F_1 : [0, T_1] \rightarrow \mathcal{P}(\mathbb{R}^n)$ and $F_2 : [0, T_2] \rightarrow \mathcal{P}(\mathbb{R}^n)$ is runtime composable if for all $x_{out} \in F_1(T_1)$, there exists a transformation Ψ_c of the state along cyclic coordinates such that $x_{out} \in \Psi_c(F_2(0))$.*

In other words, for any state x_{out} in the outlet of F_1 , one can shift F_2 along cyclic coordinates and ensure that its inlet contains x_{out} . Hence, we can guarantee that it will be possible to execute the funnel F_2 (after appropriate shifting in the cyclic coordinates) once the funnel F_1 has been executed (though the *particular* shift Ψ_c required depends on x_{out} and thus will not be known until runtime). Hence, runtime composability of funnels allows us to exploit invariances in the dynamics of the system at runtime and effectively reuse our robust motion plans in different scenarios. As a simple example, a UAV flying through a cluttered environment can reuse a funnel computed for a certain starting position by shifting the funnel so its inlet contains the UAV's current state.

Remark 3. *It is easy to see from Definition 4 that two funnels F_1 and F_2 are runtime composable if and only if:*

$$\forall x_{out} = [x_c, x_{nc}] \in F_1(T_1), \exists x_{0,c} \text{ s.t. } [x_{0,c}, x_{nc}] \in F_2(0). \quad (54)$$

This condition is simply stating that we can shift the point x_{out} along cyclic coordinates in such a way that it is contained in the inlet of F_2 . This condition in turn is equivalent to:

$$\pi_{nc}(F_1(T_1)) \subset \pi_{nc}(F_2(0)), \quad (55)$$

where as before π_{nc} denotes the projection onto the non-cyclic coordinates of the state space.




Figure 4: The figure shows two funnels that are *not* sequentially composable MI (since the inlet of funnel F_2 isn't large enough in the x_c dimension). However, since $\pi_{nc}(F_1(T_1)) \subset \pi_{nc}(F_2(0))$, they *are* runtime composable.

Figure 4 shows two funnels that are *not* sequentially composable MI. However, since $\pi_{nc}(F_1(T_1)) \subset \pi_{nc}(F_2(0))$, they *are* runtime composable.

We end our discussion on composability of funnels by noting the following relationship between the three notions of sequential composability we have discussed.

Remark 4. *The three notions of sequential composability we have discussed are related as follows:*

$$\text{Sequential composability} \implies \text{Sequential composability MI} \implies \text{Runtime composability}.$$

The first implication is immediate from Definitions 2 and 3. The second implication follows from the following reasoning:

$$\begin{aligned} F_1(T_1) &\subset \Psi_c(F_2(0)) && \text{(Sequential composability MI, ref. Definition 3)} \\ \implies \pi_{nc}(F_1(T_1)) &\subset \pi_{nc}(\Psi_c(F_2(0))) = \pi_{nc}(F_2(0)) && \text{(Runtime composability, ref. Remark 3).} \end{aligned}$$

5.4 Checking composability

Given two funnels $F_1 : [0, T_1] \rightarrow \mathcal{P}(\mathbb{R}^n)$ and $F_2 : [0, T_2] \rightarrow \mathcal{P}(\mathbb{R}^n)$ defined as $F_1(t) = \{\bar{x} \in \mathbb{R}^n \mid V_1(t, \bar{x}) \leq \rho_1(t)\}$ and $F_2(t) = \{\bar{x} \in \mathbb{R}^n \mid V_2(t, \bar{x}) \leq \rho_2(t)\}$ for polynomials V_1 and V_2 , this section describes how we can check sequential composability, sequential composability MI and runtime composability in an offline preprocessing stage.

5.4.1 Sequential compossability

Sequential compossability of (F_1, F_2) is equivalent to the following condition:

$$V_1(T_1, \bar{x}) \leq \rho_1(T_1) \implies V_2(0, \bar{x}) \leq \rho_2(0). \quad (56)$$

We can thus apply the generalized S-procedure (described in Section 3.2) and check sequential compossability using the following simple SOS program:

$$\begin{aligned} \text{Find} \quad & L(\bar{x}) \\ \text{s.t.} \quad & \rho_2(0) - V_2(0, \bar{x}) - L(\bar{x})(\rho_1(T_1) - V_1(T_1, \bar{x})) \text{ is SOS,} \\ & L(\bar{x}) \text{ is SOS.} \end{aligned} \quad (57)$$

5.4.2 Sequential composition MI

In order to check sequential compossability MI, we need to search for a shift Ψ_c along cyclic coordinates of the state such that $F_1(T_1) \subset \Psi_c(F_2(0))$. For the important special case in which the sets $F_1(T_1)$ and $F_2(0)$ are ellipsoids (corresponding to $V_1(T_1, \bar{x})$ and $V_2(0, \bar{x})$ being quadratic² in \bar{x}), we can cast this search as a semidefinite program (SDP). Suppose the sets $F_1(T_1)$ and $F_2(0)$ are given by:

$$F_1(T_1) = \{x \in \mathbb{R}^n \mid (x - x_1(T_1))^T S_1(x - x_1(T_1)) \leq 1\} \quad (58)$$

$$F_2(0) = \{x \in \mathbb{R}^n \mid (x - x_2(0))^T S_2(x - x_2(0)) \leq 1\}, \quad (59)$$

where $x_1 : [0, T_1] \rightarrow \mathbb{R}^n$ and $x_2 : [0, T_2] \rightarrow \mathbb{R}^n$ are the nominal trajectories around which the funnels were computed and S_1 and S_2 are positive definite matrices. We would like to search for a shift $\Delta_c \in \mathbb{R}^n$ (where the components of Δ_c corresponding to the non-cyclic coordinates are set to zero) such that the ellipsoid:

$$\Psi_c(F_2(0)) = \{x \in \mathbb{R}^n \mid [x - (x_2(0) + \Delta_c)]^T S_2[x - (x_2(0) + \Delta_c)] \leq 1\} \quad (60)$$

is contained within the ellipsoid $F_1(T_1)$. The first step in obtaining the desired SDP is to note that the set $\Psi_c(F_2(0))$ can be represented equivalently as the image of the unit ball under an affine map:

$$\Psi_c(F_2(0)) = \{Bu + x_2(0) + \Delta_c \mid \|u\|_2 \leq 1\}, \quad (61)$$

where $B = \text{chol}(S_2)^{-1}$. Here, $\text{chol}(S_2)$ is the Cholesky factorization of S_2 (guaranteed to exist and be invertible since S_2 is positive definite). Such a representation is a standard trick in semidefinite programming (see for example [Boyd and Vandenberghe, 2004, Section 8.4.2]).

Introducing the notation $b := -S_1 x_1(T_1)$ and $c := x_1(T_1)^T S_1 x_1(T_1) - 1$, the condition that the ellipsoid $F_1(T_1)$ is a subset of the ellipsoid $\Psi_c(F_2(0))$ is then *equivalent* to being able to find a scalar $\lambda > 0$ such that the following matrix semidefiniteness constraint holds ([Boyd and Vandenberghe, 2004, Section 8.4.2]):

$$\begin{bmatrix} -\lambda - c + b^T S_1^{-1} b & 0_{1 \times n} & (x_2(0) + \Delta_c + S_1^{-1} b)^T \\ 0_{n \times 1} & \lambda I_{n \times n} & B \\ x_2(0) + \Delta_c + S_1^{-1} b & B & S_1^{-1} \end{bmatrix} \succeq 0. \quad (62)$$

²Note that the restriction to quadratic $V_1(T_1, \bar{x})$ and $V_2(0, \bar{x})$ is a relatively mild one. We are not imposing any conditions on the degree of V_1 and V_2 at times other than the endpoints.

Here, the matrices I and 0 represent the identity matrix and the all-zeros matrix respectively. Since this semidefiniteness condition is linear in λ and Δ_c , the problem of searching for these decision variables subject to $\lambda > 0$ and (62) is a SDP. This SDP will be feasible if and only if F_1 and F_2 are sequentially composable MI.

The problem of checking sequential composable MI in the more general non-ellipsoidal case is not directly amenable to such a SDP based formulation. However, if we *fix* Ψ_c , then the problem of checking sequential composable MI reduces to the problem of checking sequential composable. Thus, we can use the SOS program (57) to verify if a *given* Ψ_c yields the desired containment constraint $F_1(T_1) \subset \Psi_c(F_2(0))$. One natural choice is to set Ψ_c such that $\pi_c(\Psi_c(x_2(0))) = \pi_c(x_1(T_1))$, where π_c is the projection of the state onto the cyclic coordinates. Intuitively, this corresponds to shifting the funnel F_2 so that the start of its nominal trajectory is lined up along cyclic coordinates with the end of the nominal trajectory of F_1 . If the SOS program is infeasible, a local search around this Ψ_c could yield the desired shift.

5.4.3 Runtime composable

In order to check runtime composable of F_1 and F_2 , we need to check the inclusion $\pi_{nc}(F_1(T_1)) \subset \pi_{nc}(F_2(0))$. For the important special case where the sets $F_1(T_1)$ and $F_2(0)$ are ellipsoids, we can compute the projections exactly. This is because the projection of an ellipsoid onto a linear subspace is also an ellipsoid (see equation (45) for the exact formula for the projected ellipsoid). Checking if a given ellipsoid contains another is a straightforward application of semidefinite programming [Boyd and Vandenberghe, 2004, Example B.1, Appendix B].

For the more general case, one might hope for a SOS programming based condition for checking $\pi_{nc}(F_1(T_1)) \subset \pi_{nc}(F_2(0))$. However, the existential quantifier inherent in the projection (see the equivalent condition (54)) makes it challenging to formulate such SOS conditions. Nevertheless, there exist general purpose tools such as *quantifier elimination* [Collins and Hong, 1998] for checking quantified polynomial formulas such as (54). While the worst-case complexity of doing general purpose quantifier elimination is poor, software packages such as QEPCAD [Brown, 2003] (or dReal [Gao et al., 2013], which is based on a different theoretical framework) can often work well in practice for specialized problems. We note however that in the examples considered in Sections 7 and 8 we will only be using funnels with ellipsoidal inlets and outlets and thus will not be concerned with this complexity.

5.5 Funnel library

A simple but useful generalization of the notions of composable introduced above can be obtained by checking the associated containment conditions at a given time τ_1 rather than at time T_1 . For example, we will say that the ordered pair of funnels (F_1, F_2) is sequentially composable at time τ_1 if $F_1(\tau_1) \subset F_2(0)$. We will use similar terminology for the other notions of composable. Given a collection \mathcal{F} of funnels associated with a dynamical system, we will associate a directed graph $\mathcal{G}(\mathcal{F})$ whose vertices correspond to funnels $F \in \mathcal{F}$. Two vertices corresponding to funnels F_i and F_j are connected by a directed edge (F_i, F_j) if and only if the ordered pair (F_i, F_j) is runtime composable at some specified time τ_i . We will sometimes refer to τ_i as the execution time of funnel F_i .

Definition 5. A funnel library FL associated with a given dynamical system is a tuple $FL = (\mathcal{F}, \mathcal{G}(\mathcal{F}), \mathcal{C}, \{\tau_i\})$, where \mathcal{F} is a set of funnels for the dynamical system, $\mathcal{G}(\mathcal{F})$ is the directed graph

representing which funnels are runtime composable, \mathcal{C} is the set of feedback controllers associated with the funnels in \mathcal{F} , and $\{\tau_i\}$ is the set of execution times.

Note that while we do not impose restrictions such as connectedness or strong connectedness on the graph $\mathcal{G}(\mathcal{F})$, it may be useful to impose such conditions based on the task at hand. For example, for tasks which require continuous operation (such as a UAV navigating indefinitely through a forest or a factory arm continuously placing objects onto a conveyor belt), we should require that $\mathcal{G}(\mathcal{F})$ be strongly connected (since if $\mathcal{G}(\mathcal{F})$ is not strongly connected, we are ruling out the possibility of executing certain funnels in the future). On a similar note, properties of the graph such as its *diameter* or *girth* may be related to the efficiency with which a certain task can be accomplished. Identifying which graph theoretic properties should be imposed on $\mathcal{G}(\mathcal{F})$ for different tasks is an interesting research avenue, but we do not pursue this in the present work.

6 Real-time Planning with Funnels

Given a funnel library $\text{FL} = (\mathcal{F}, \mathcal{G}(\mathcal{F}), \mathcal{C}, \{\tau_i\})$ computed offline, we can proceed to use it for robust real-time planning in previously unseen environments. The robot's task specification may be in terms of a goal region that must be reached (as in the case of a manipulator arm grasping an object), or in terms of a nominal direction the robot should move in while avoiding obstacles (as in the case of a UAV flying through a forest or a legged robot walking over rough terrain). For the sake of concreteness, we adopt the latter task specification although one can adapt the contents of this section to the former specification. We further assume that the robot is provided with regions in the configuration space that obstacles are guaranteed to lie in and that the robot's sensors only provide this information up to a finite spatial horizon around the robot. Our task is to choose funnels from our library in a way that avoids obstacles while moving forward in the nominal direction.

The key step in our real-time planning approach is the selection of a funnel from the funnel library that doesn't intersect any obstacles in the environment. This selection process is sketched in the ReplanFunnel algorithm (Algorithm 3). Given the current state x of the robot and the locations and geometry of obstacles \mathcal{O} in the environment, the algorithm searches through the funnels in the library that are runtime composable with the previous funnel that was executed. We assume that the funnels are (totally) ordered in a preference list. As an example, this ordering can be based on likely progress towards the goal or by aggressiveness of the maneuvers (less aggressive maneuvers are given preference) for a UAV. For each funnel F_i , the algorithm tries to find a shift Ψ_c along cyclic coordinates of the system such that the shifted funnel $\Psi_c(F_i)$ satisfies two properties: (i) the current state x is contained in the inlet of the shifted funnel, (ii) the projection of the shifted funnel onto the coordinates of the state space corresponding to the configuration space³ doesn't intersect any obstacles in \mathcal{O} . If we are able to find such a shift, we are *guaranteed* that the system will remain collision-free when the funnel is executed despite the uncertainties and disturbances that the system is subjected to.

The simplest way to try to choose Ψ_c is to set it such that the nominal trajectory x_i associated with the funnel lines up with the current state in the cyclic coordinates (i.e., using the notation of Section 5, $\pi_c(\Psi_c(x_i(0))) = \pi_c(x)$). Intuitively, this corresponds to shifting the funnel so that it is executed from the current location of the robot. We can then use standard collision-checking

³Note that these projections can be computed in the offline computation stage.

libraries such as the Bullet Collision Detection & Physics Library [Coumans, 2014] to check if $\Psi_c(F_i)$ (projected onto the configuration space) intersects any obstacles. We will discuss more sophisticated ways of finding Ψ_c in Section 6.1.

If the search for a collision-free funnel in Algorithm 3 doesn't succeed, we assume that there is a failsafe maneuver that can be executed to keep the robot safe. For a ground vehicle, this could entail coming to a stop while for a quadrotor or fixed-wing airplane this may involve transitioning to a hover or propellor-hang mode. In certain cases, it is possible to derive geometric conditions on the environment that the robot will operate in (e.g., constraints on obstacle size and gaps between obstacles) that *guarantee* that a collision-free funnel will always be found by Algorithm 3 if the environment satisfies these conditions. We will see an example of this in Section 7.2 for a quadrotor system navigating through a forest of polytopic obstacles.

```

1: Inputs:  $x$  (current state of system),  $\mathcal{O}$  (reported obstacles in environment), previousFunnel
   (previous funnel that was executed)
2: for  $i = 1, \dots, \#\mathcal{F}$  such that (previousFunnel,  $F_i \in \mathcal{G}(\mathcal{F})$ ) do
3:   success  $\Leftarrow$  Find a shift  $\Psi_c$  along cyclic coordinates such that  $x$  is contained in the inlet of
    $\Psi_c(F_i)$  and  $\Psi_c(F_i)$  is collision-free w.r.t.  $\mathcal{O}$ 
4:   if success then
5:     return  $\Psi_c(F_i)$ 
6:   end if
7: end for
8: return  $F_{failsafe}$ 
```

Algorithm 3: ReplanFunnel

Algorithm 4 provides a sketch of the real-time planning and control loop, which applies the ReplanFunnel algorithm in a receding-horizon manner. At every control cycle, the robot's sensors provide it with a state estimate and report the locations and geometry of the set of obstacles \mathcal{O} in the sensor horizon. The algorithm triggers a replanning of funnels if any of the following three criteria are met: (i) if the system has executed the current funnel F_i for the associated execution time τ_i , (ii) if the current state of the system is no longer in the funnel being executed, or (iii) if the current funnel being executed is no longer collision-free. In principle, (ii) should not happen. However, in practice this can happen if the system received a disturbance that was larger than the ones taken into account for the funnel computations. Option (iii) can happen if the robot's sensors report new obstacles that were previously unseen.

6.1 Shifting funnels at runtime

The main step in Algorithm 3 is the search for a shift Ψ_c along cyclic coordinates such that the shifted funnel $\Psi_c(F_i)$ will be collision-free with respect to the obstacles in \mathcal{O} while containing the current state x in its inlet:

$$\text{Find} \quad \Psi_c \quad (63)$$

$$\text{s.t.} \quad x \in \Psi_c(F_i(0)), \quad (64)$$

$$\pi_{conf}(\Psi_c(F_i)) \cap \mathcal{O} = \emptyset.$$

```

1:  $x \leftarrow$  Initialize current state of the robot
2:  $\mathcal{O} \leftarrow$  Initialize obstacles in sensor horizon
3:  $F_i \in \mathcal{F} \leftarrow$  Initialize current planned funnel
4: for  $t = 0, \dots$  do
5:    $x \leftarrow$  Update current state of robot
6:    $\mathcal{O} \leftarrow$  Update obstacles in sensor horizon
7:   replan  $\leftarrow$  Check if we have finished executing current funnel  $F_i$  for the associated execution
   time  $\tau_i$ 
8:   insideFunnel  $\leftarrow$  Check if current state is still inside the current funnel  $F_i$  being executed
9:   collisionFree  $\leftarrow$  Check if current funnel  $F_i$  is still collision-free with  $\mathcal{O}$ 
10:  if replan or  $\neg$ insideFunnel or  $\neg$ collisionFree then
11:     $F_i \leftarrow$  ReplanFunnels( $x, \mathcal{O}, F_i$ )
12:  else
13:    apply feedback control input  $u$  associated with current funnel  $F_i$ 
14:  end if
15: end for

```

Algorithm 4: Real-time planning loop

Here, π_{conf} is the projection onto the configuration space variables of the state space. This optimization problem is non-convex in general since the free-space of the environment is non-convex. However, the number of decision variables is very small. In particular, if we parameterize the shift Ψ_c by a vector $\Delta_c \in \mathbb{R}^n$ (where the coordinates of Δ_c corresponding to the non-cyclic coordinates are set to zero) such that $\Psi_c(x) = x + \Delta_c$, the number of decision variables is equal to the number of cyclic coordinates of the system.

We can thus apply general-purpose nonlinear optimization tools such as gradient-based methods to solve this problem. In particular, the first constraint is equivalent to checking that the value of the Lyapunov function $V(0, \bar{x}) = V(0, x - (x_i(0) + \Delta_c))$ at time 0 is less than or equal to $\rho(0)$ (recall that the funnel is described as the $\rho(t)$ sub-level set of the function $V(t, \bar{x})$). As in Section 4, x_i here is the nominal trajectory corresponding to the funnel F_i . The second constraint can be evaluated using off-the-shelf collision-checking libraries.

Despite the small number of decision variables, for some applications general-purpose nonlinear optimization may be too slow to run in real-time. For the important special case where the cyclic coordinates form a subset of the configuration space variables (e.g., for a UAV), we propose another approach that allows us to search over a restricted family of shifts Ψ_c but has the advantage of being posed using *convex* quadratic constraints. The first observation is that the containment constraint (64) is a convex quadratic constraint for the case where the inlet of the funnel is an ellipsoid (i.e., the function $V(0, \bar{x})$ is a positive definite quadratic). If the inlet is not an ellipsoid, we can find an ellipsoidal inner approximation by solving a simple SOS program offline.

Next, we seek to find a set of convex constraints that will guarantee non-collision of the funnel. We will assume that $\pi_{conf}(F_i)$ is represented as a union of convex segments \mathcal{S}_k and that the obstacles in the environment are also convex (we assume that non-convex obstacles have been decomposed into convex segments). For each funnel segment, we can find collision normals n_{jk} and collision distances d_{jk} to each obstacle $o_j \in \mathcal{O}$ (these are easily extracted from a collision-checking software). Figure 5 provides an illustration of this for a single convex segment \mathcal{S}_1 . The two regions colored in red are obstacles. The collision normals are n_{11} and n_{21} . By definition, collision normals




Figure 5: The two red regions are obstacles. The green region represents the set of allowable shifts of the funnel segment that satisfy the linear constraints given by the collision normals and collision distances.

provide us with constraints such that any shift Δ_c of the segment S_1 satisfying the linear constraint $n_{jk}^T \Delta_c < d_{jk}$ will be collision-free. The region corresponding to this is shaded in green in the figure.

The constraints described above (containment of the current state in an ellipsoid and the linear constraints given by the collision normals) form a special case of convex Quadratically Constrained Quadratic Programs (QCQPs), for which there exist very mature software packages. For our examples in Section 7, we use the FORCES Pro package [Domahidi and Jerez, 2014]. The package generates solver code tailored to the specific optimization problem at hand and is faster in our experience than using general-purpose convex QCQP solvers.

6.2 Global Planning

Algorithm 4 employed a receding horizon strategy for real-time planning. For tasks such as robot navigation through previously unmapped environments, such a replanning-based strategy is unavoidable since the robot's sensors will only report obstacles in the environment in some finite sensor horizon around the robot. Thus, the robot will need to replan as it makes progress through the environment and more obstacles are reported by its sensors. However, in certain scenarios where the robot has access to a larger portion of the environment, a planning strategy that is more *global* in nature may be appropriate. In general, the funnel primitives provide a discrete action space which can be searched by any heuristic planner - the primary considerations here are the

additional constraint of containment of the current state in the inlet and the moderately more significant cost of collision checking. Here we briefly mention a few ways in which global planning may be achieved.

Grid-based planners: One can build on existing trajectory library-based planning approaches to quickly find collision-free sequences of funnels that minimize a certain performance criterion (e.g., distance traversed by the nominal trajectories in the sequence). For example, [Stolle and Atkeson, 2006] uses a variant of the grid-based planning algorithm A^* to search through possible sequences of trajectories in a library. The additional cost of extending such an approach for planning with funnels is the cost that comes from pruning sequences of funnels that are in collision with obstacles in the environment.

Maneuver automata: The Maneuver Automaton [Frazzoli et al., 2005] provides an alternative trajectory library-based approach that exploits continuous symmetries (such as shift invariance for ground vehicles or UAVs) to achieve efficient planning. The approach relies on having a number of “trim trajectories” and maneuvers that transition between them. One can then efficiently plan sequences of trims and maneuvers that start and end at prescribed points in space by solving a set of equations that has the same structure as an inverse kinematics problem. In order to extend this approach for real-time planning with funnels, one needs to (potentially approximately) solve the inverse kinematics problem subject to the constraint that the funnel sequence is collision-free.

Planning backwards: The real-time planning approaches we have discussed so far all plan forward in time. However, in cases where there is a well-defined goal set that one needs to reach, it may be more efficient to plan backwards. This kind of planning is a direct analog of the *preimage backchaining* approach [Lozano-Perez et al., 1984] in the motion planning literature. One way to perform this backwards search would be to employ a randomized strategy similar to the Rapidly-exploring Randomized Tree (RRT). In particular, we can grow a tree of funnels backwards from the goal set by using the funnels in our funnel library as primitives for the extension operator for the RRT. One can again exploit invariances in the dynamics when performing this extension operation by randomly sampling shifts/rotations of funnels (while still maintaining sequential compossibility and non-collision of funnels). The termination criterion for this RRT-like algorithm is the containment of the current state of the robot in the tree of funnels. This is very similar to the LQR-Trees approach [Tedrake et al., 2010], with the following differences: (i) one would only use a pre-computed library of funnels in the tree rather than computing new funnels (which would be computationally infeasible for real-time implementation), and (ii) the extension operator in the tree would take into account non-collision of funnels with obstacles in the environment.

7 Examples

In this section we apply the techniques presented in this paper on two simulation examples. We will consider a hardware example in Section 8. The computations in this section were performed on a 3.4 GHz desktop computer with 16 GB RAM and 4 cores.




Figure 6: Illustration of the ground vehicle model.

7.1 Ground Vehicle Model

As our first example we consider a ground vehicle model based on the Dubins car [Dubins, 1957] navigating an environment of polytopic obstacles. A pictorial depiction of the model is provided in Figure 6. The vehicle is constrained to move at a fixed forward speed and can control the second derivative of its yaw angle ψ . We introduce uncertainty into the model by assuming that the speed of the vehicle is only known to be within a bounded range and is potentially time-varying. The full non-linear dynamics of the system are then given by:

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ \psi \\ \dot{\psi} \end{bmatrix}, \quad \dot{\mathbf{x}} = \begin{bmatrix} -v(t) \sin \psi \\ v(t) \cos \psi \\ \dot{\psi} \\ u \end{bmatrix} \quad (65)$$

with the speed of the plane $v(t) \in [9.0, 11.0]$ m/s. The control input is bounded in the range $[-1000, 1000]$ rad/s².

The trajectory library, \mathcal{T} , computed for the ground vehicle consists of 20 trajectories and is shown in Figure 7(a). The trajectories $x_i(t) : [0, T_i] \mapsto \mathbb{R}^4$ and the corresponding nominal open-loop control inputs were obtained via the direct collocation trajectory optimization method [Betts, 2001] for the vehicle dynamics in (65) with $v(t) = 10$ m/s. The initial state $x_i(0)$ was constrained to be $[0, 0, 0, 0]$ and the final state $x_i(T_i)$ was varied in the x direction while keeping the y component fixed. We locally minimized a cost of the form:

$$J = \int_0^{T_i} [1 + u_0(t)^T R(t) u_0(t)] dt$$

where R is a positive-definite matrix. We constrained the nominal control input to be in the range $[-500, 500]$ rad/s² to ensure that feedback controllers computed around the nominal trajectories do not immediately saturate.

For each $x_i(t)$ in \mathcal{T} we obtain controllers and funnels using the method described in Section 4. In order to obtain polynomial dynamics, we computed a (time-varying) degree 3 Taylor expansion




Figure 7: The plot on the left shows the trajectory library for the ground vehicle model. The plot on the right shows a selection of funnels from the funnel library projected onto the $x - y$ plane.

of the dynamics of the system around the nominal trajectory. We note that with the right change of coordinates, one can express the dynamics of this system directly as a polynomial. In particular, we can introduce new indeterminates s and c for $\sin(\psi)$ and $\cos(\psi)$, and impose the constraint that $s^2 + c^2 = 1$ (this equality constraint is easily imposed in the sums-of-squares programming framework). However, this increases the dimensionality of the state space and in practice we find that the time-varying Taylor approximation accurately captures the nonlinearities of the system.

The approach from Section 4.3.2 along with the time-sampled approximation described in Section 4.2 (with 15 time samples) was used to synthesize a (time-varying) linear feedback controller around each trajectory. The methods described in Section 4.3.1 and 4.3.3 were used to take into account the parametric uncertainty and input saturations that the system is subject to. As described in Section 4.4.2, we used a time-varying LQR controller to initialize the funnel computations. The computation time for each funnel was approximately 5-10 minutes. A subset of the funnels is shown in Figure 7(b). Note that the four-dimensional funnels have been projected down to the $x - y$ dimensions for the purpose of visualization. The directed graph $\mathcal{G}(\mathcal{F})$ that encodes real-time composability between funnels (ref. Section 5.5) is fully connected.

The resulting funnel library was employed by Algorithm 4 for planning in real-time through environments with randomly placed obstacles. Figure 8 shows the funnels that were executed in order to traverse a representative environment. The obstacle positions were randomly generated from a spatial Poisson process (with a density/rate parameter of 0.6 obstacles per m^2). Two further ‘‘barrier’’ obstacles were placed on the sides of the environment to prevent the vehicle from leaving the region containing obstacles. The planner was provided with a sensor horizon of 3m in the y direction (forward) and ± 2 m in the x direction (side-to-side) relative to the position of the vehicle. Only obstacles in this sensor window were reported to the planner at every instant in time. The execution times τ_i for each funnel were set such that replanning occurred once 80% of the funnel was executed. The parametric uncertainty in the speed of the vehicle was taken into account in our simulation by randomly choosing a speed $v(t) \in \{9.0, 11.0\}$ m/s after every execution time period has elapsed. The real-time planner employs the QCQP-based algorithm from Section 6.1




Figure 8: This plot shows the funnels that were executed in order for the ground vehicle model to traverse a randomly generated obstacle environment using the funnel library based real-time planning approach.

for shifting funnels in the $x - y$ directions (we did not exploit invariances in the yaw dimension of the state space in order to ensure that the vehicle moves in the forward direction and doesn't veer off to the sides). We use the FORCES Pro solver [Domahidi and Jerez, 2014] for our QCQP problems. This resulted in our implementation of the real-time planner running at approximately 40 – 50 Hz.

We performed extensive simulation experiments to compare our funnel library based robust planning approach with a more traditional trajectory library based method. In order to facilitate a meaningful comparison, we used the underlying trajectory library corresponding to our funnel library. The trajectory-based planner employs essentially the same outer-loop as our funnel-based approach (Algorithm 4). The key difference is that the planner chooses which maneuver to execute by evaluating which trajectory has the maximal clearance from the obstacles (as measured by Euclidean distance):

$$\max_i \min_{t,j} \text{dist}(x_i(t), o_j) \quad (66)$$

where $x_i(t) \in \mathcal{T}$ is a trajectory in the library and $o_j \in \mathcal{O}$ is an obstacle in the environment. Once a maneuver is chosen based on this metric, a time-varying LQR feedback controller computed along this trajectory is applied (the same LQR controller that is used to initialize the funnel computations).

To compare the two planners, we generated 100 obstacle environments randomly as described previously. For each environment, we ran the different planning algorithms until there was a collision of the vehicle with an obstacle. The distance in the y direction at the time of collision was recorded for each run. Figure 9 compares the performance of the different approaches. For each distance on the x-axis of the plot the height of the bar indicates the number of runs for which the vehicle traveled beyond this distance. As is evident from the plot, our funnel library based approach provides a significant advantage over the trajectory-based method.




Figure 9: A bar plot comparing the performance of the funnel library approach with one based on a trajectory library.

This advantage can be partially understood by considering a specific example. In particular, Figure 10 demonstrates the utility of explicitly taking into account uncertainty during the planning process. There are two obstacles in front of the vehicle. The two options available to the plane are to fly straight in between the obstacles or to maneuver aggressively to the right and attempt to go around them. If the motion planner didn't take uncertainty into account and simply chose to maximize the clearance in terms of Euclidean distance from the nominal trajectory to the obstacles

(see equation (66)), it would choose the trajectory that goes right around the obstacles. However, taking the funnels into account leads to a different decision: going straight in between the obstacles is guaranteed to be safe even though the distance to the obstacles is smaller. In contrast, the maneuver that avoids going in between obstacles is less robust to uncertainty and could lead to a collision. The utility of safety guarantees in the form of funnels is especially important when the margins for error are small and making the wrong decision can lead to disastrous consequences.




Figure 10: This figure shows the utility of explicitly taking uncertainty into account while planning. The intuitively more risky strategy of flying in between two closely spaced obstacles is guaranteed to be safe, while the path that avoids going in between obstacles is less robust to uncertainty and could lead to a collision.

7.2 Quadrotor Model

The next example we consider is a model of a quadrotor system navigating through a forest of polygonal obstacles. A visualization of the system is provided in Figure 11. The goal of this example is to demonstrate that we can derive simple geometric conditions on the environment that guarantee collision-free flight. In other words if the environment satisfies these conditions, the Algorithm 3 presented in Section 6 will always succeed in finding a collision-free funnel from the library and the quadrotor will fly forever through the environment with no collisions.

The quadrotor model has a 12 dimensional state space consisting of the x-y-z position of the centre of mass, the roll-pitch-yaw of the body, and the time derivatives of these configuration space




Figure 11: Visualization of the quadrotor system.

variables. The dynamics model we use is identical to the one presented in [Mellinger et al., 2010] with an additional uncertainty term in the form of a “cross wind”. We modeled this with a bounded uncertainty term on the acceleration of the x position: $\ddot{x} = \ddot{x}_{nominal} + \Delta$, with $\Delta \in [-0.1, 0.1] \text{ m/s}^2$.

Figure 12(a) plots the trajectory library we use. The funnel library consists of 20 maneuvers and was created in a manner similar to the ground vehicle example. In particular, the trajectories $x_i(t) : [0, T_i] \mapsto \mathbb{R}^{12}$ and the corresponding nominal open-loop control inputs were obtained via the direct collocation trajectory optimization method [Betts, 2001]. The initial state $x_i(0)$ was constrained to have a forwards speed of 2 m/s (i.e., $x_i(0) = [0, 0, 0, 0, 0, 0, 0, 0, 2, 0, 0, 0, 0]$) and the final state $x_i(T_i)$ was varied in the x direction while keeping all other components of the state fixed. We locally minimized a cost of the form:

$$J = \int_0^{T_i} [1 + u_0(t)^T R(t) u_0(t)] dt$$

where R is a positive-definite matrix. In addition to the 20 maneuvers, the library also consists of a “trim” trajectory corresponding to the quadrotor flying forward at a constant speed of 2 m/s.

For each $x_i(t)$ in \mathcal{T} we obtain controllers and funnels using the method described in Section 4. We obtained time-varying Taylor expansions of degree 3 computed around the nominal trajectory. The approach from Section 4.3.2 along with the time-sampled approximation described in Section 4.2 (with 15 time samples) was used to synthesize a (time-varying) linear feedback controller around each trajectory. The methods described in Section 4.3.1 were used to take into account the parametric uncertainty that the system is subject to. The computation time for each funnel was approximately 20-25 minutes. The directed graph $\mathcal{G}(\mathcal{F})$ that encodes real-time composability between funnels (ref. Section 5.5) is fully connected. Moreover, the funnel corresponding to the trim trajectory is sequentially composable modulo invariances (ref. Section 5.2) with the other maneuvers in the library, allowing us to apply the trim trajectory before or after any of the maneuvers. A subset of the funnels is shown in Figure 12(b). Note that the twelve-dimensional funnels have been projected down to the $x - y - z$ dimensions for the purpose of visualization.

Given this funnel library \mathcal{F} , we will show how one can derive simple geometric conditions on the environment that guarantee that a collision-free funnel will always be found during real-




Figure 12: The plot on the left shows the trajectory library for the quadrotor model. The plot on the right shows a selection of funnels from the funnel library projected onto the $x - y - z$ space.

time planning. In order to simplify the analysis, we will assume that the quadrotor is navigating through a 2.5D environment (i.e., the obstacles in the environment are 2D polygons extruded in the z -direction). But, we note that our analysis can be extended to fully three dimensional environments. We will treat the geometry of the quadrotor as a sphere of radius r_{quad} .

Since we are only considering 2.5D environments, it is sufficient to project the environment and funnels down to the $x - y$ plane. Consider (axis aligned) bounding boxes for the 2D obstacles (see Figure 13). Denote the center of the bounding box for obstacle o_i as $(o_{i,x}, o_{i,y})$ for $i \in \{1, \dots, N_{obs}\}$ (N_{obs} is the number of obstacles in the robot's sensor horizon). Suppose without loss of generality that the quadrotor's x - y - z position is $(0, 0, 0)$. We will first derive conditions on the distances $o_{i,y}$, $d_{x,ij} := |o_{i,x} - o_{j,x}|$, and sizes $r_{x,i}$, $r_{y,i}$ (ref. Figure 13) such that there exists a collision-free funnel in \mathcal{F} *assuming* the quadrotor's x - y - z position is $(0, 0, 0)$. We will then extend this analysis to derive conditions that guarantee that a collision-free funnel will always be found as the quadrotor applies Algorithm 3 in a receding horizon manner to fly continuously through the environment.

For each funnel $F_i \in \mathcal{F}$ compute bounding boxes for the inlet and outlet of the funnel (see Figure 14(b)). We will refer to the dimensions of these bounding boxes as depicted in the figure as $f_{i,y}^{out}, f_{i,x}^{out}, f_{i,y}^{in}, f_{i,x}^{in}$. Define

$$f_x := \max(f_{1,x}^{out}, \dots, f_{N,x}^{out}, f_{1,x}^{in}, \dots, f_{N,x}^{in}) + r_{quad},$$

where N is the number of funnels in the library (21 in our case). Similarly, define:

$$f_y := \max(f_{1,y}^{out}, \dots, f_{N,y}^{out}, f_{1,y}^{in}, \dots, f_{N,y}^{in}) + r_{quad}.$$

Next, define δ as shown in Figure 14(a) as the distance (in the x -dimension) of the endpoints of trajectories in \mathcal{T} . Denote by Δ_x the distance (in the x -dimension) between the endpoints of the most aggressive trajectories in \mathcal{T} and let Δ_y be the distance in the y -direction that each trajectory in the library covers (see Figure 14(a)).

Define $r_x := \max(r_{1,x}, \dots, r_{N_{obs},x})$, $r_y := \max(r_{1,y}, \dots, r_{N_{obs},y})$. Then consider the following conditions on the obstacles:




Figure 13: Obstacle bounding boxes and quadrotor radius.

$$o_{i,y} > \Delta_y, \forall i \in \{1, \dots, N_{obs}\} \quad (67)$$

$$r_y < f_y \quad (68)$$

$$r_x < \Delta_x - 2f_x \quad (69)$$

$$d_{x,ij} > 2f_x + \delta, \forall i, j \in \{1, \dots, N_{obs}\}, i \neq j. \quad (70)$$

It is straightforward to see that these conditions imply the existence of a funnel in \mathcal{F} that is collision-free when executed from the quadrotor's location at $(0, 0, 0)$ (i.e., Algorithm 3 will succeed assuming that the full 12 dimensional state of the quadrotor is contained within the inlet of all the funnels in \mathcal{F}). In particular, the condition (67) ensures that the funnels are not too close to the quadrotor's starting location. Condition (68) prevents an obstacle from extending so far in the y -direction that it collides with the middle segments of a funnel. Condition (69) ensures that obstacles do not extend so far in the x -direction that there is no funnel in the library that goes around the obstacle (e.g., imagine a wall in front of the quadrotor). And finally (70) ensures that the gap in the x -direction between any two obstacles is large enough for *some* funnel in \mathcal{F} to fit through. Then conditions (69) and (70) ensure that this gap occurs in a portion of the space that is reachable by one of the funnels in the library (and not in some far off location in the x -direction where a funnel in \mathcal{F} does not extend to).

We can now extend this analysis to the scenario in which the quadrotor is flying along continuously through the environment by applying Algorithm 3 in a receding horizon manner. We will assume that the sensor horizon of the quadrotor (i.e., the range in which obstacles are reported) is greater than Δ_y in the y -direction and greater than $\Delta_x/2$ in the x -direction. We then need to




Figure 14: Notation for geometry of trajectories and funnels.




Figure 15: The plot shows the quadrotor maneuvering through a forest of polytopic obstacles in a collision-free manner. The environment satisfies simple geometric conditions that allow us to guarantee collision-free flight forever. Note that the visualized funnels have been inflated to take into account the size of the quadrotor.

ensure that when the quadrotor finishes executing a funnel, it does not find itself in a position where there are obstacles that are very close to it. In order to do this, we will first replace the conditions (67) and (70) with the following condition on the separation between obstacles:

$$d_{y,ij} := |o_{i,y} - o_{j,y}| > \Delta_y + f_y \quad (71)$$

OR

$$d_{x,ij} > 2f_x + \delta. \quad (72)$$

This condition ensures that obstacles are separated enough in *either* the x or y directions. However, this is still not enough to prevent cases where the quadrotor executes a funnel and finds itself too close to an obstacle. In particular, suppose that the quadrotor starts off at location $(0, 0, 0)$ and that there is a single obstacle in the environment located at $(0, 1.1\Delta_y + f_y)$ (the size of the obstacle will not matter in this example). When the quadrotor applies Algorithm 3 at location $(0, 0, 0)$ to find a funnel, all the funnels in the library will be collision-free. Let us suppose that the algorithm chooses the funnel corresponding to the quadrotor flying straight and ends up in location $(0, \Delta_y)$. At this point, the obstacle is too close to the quadrotor and a collision-free funnel will not be found. To prevent cases such as this, we can use the trim maneuver in \mathcal{F} corresponding to the quadrotor flying straight in order to “pad” the distance to the obstacles. In particular, we can apply the trim maneuver until the quadrotor is at distance Δ_y from the obstacle (in the y -direction) and *then* use Algorithm 3 to replan. This will guarantee that Algorithm 3 will succeed and thus our analysis is complete.

Figure 15 shows an example of the quadrotor system navigating through an environment that

satisfies the geometric conditions derived above. The figure shows the sequence of funnels executed by the quadrotor to traverse the environment. Note that the funnels depicted in the plot have been inflated by r_{quad} in order to take into account the physical extent of the quadrotor.

We end this section by noting that our goal here has been to demonstrate the possibility of imposing geometric conditions on the environment that guarantee collision-free flight. Going forwards, our analysis may be varied or tightened in many ways. For example, we can extend it in a straightforward manner to fully 3D environments. Further, we have assumed that the quadrotor is only planning a single funnel at a time (in contrast to sequences of funnels). Planning sequences of funnels may help loosen some of the restrictions on the environment that we have made in our analysis.

8 Hardware experiments on a fixed-wing airplane

In this section, we will validate the key components of the approach presented in this paper on a small fixed-wing airplane performing a challenging obstacle avoidance task. The goal of these hardware experiments is to answer the following important practical questions:

- Can we obtain models of a real-world challenging nonlinear dynamical system that are accurate enough to compute funnels that are valid in reality?
- Can we implement the real-time planning algorithm described in Section 6 to operate at the required rate given realistic computational constraints?
- Can we demonstrate our planning algorithm on a realistic and challenging obstacle avoidance task?

8.1 Hardware platform


The hardware platform chosen for these experiments is the SBach RC airplane manufactured by E-flite shown in Figure 16(a). The airplane is very light (76.6 g) and highly maneuverable, thus allowing for dramatic obstacle avoidance maneuvers in a tight space. The control inputs to the SBach are raw servo commands to the control surfaces (ailerons, rudder, elevator) and a raw throttle setting. These commands are sent through a modified 2.4 GHz RC transmitter at an update rate of 50 Hz.

8.2 Task and experimental setup

The experimental setup is shown in Figure 16(b). The airplane is launched from a simple rubber-band powered launch mechanism at approximately 4 – 5 m/s. The goal is to traverse the length of the room while avoiding the obstacles placed in the experimental arena. The airplane’s planner is *not* informed where the obstacles are beforehand; rather, the obstacle positions and geometry are reported to the planner only once the airplane clears the launcher. This simulates the receding-horizon nature of realistic obstacle avoidance tasks where the obstacle positions are not known beforehand and planning decisions must be taken in real-time. The experiments are performed in a Vicon motion capture arena that reports the airplane and obstacle poses at 120 Hz. All the online computation is performed on an off-board computer with four Intel i7 2.9GHz processors and 16 GB RAM.



(a) The E-flite SBach RC airplane used for hardware experiments.



(b) Experimental setup and coordinate system.

Figure 16: Airplane hardware and experimental setup.

8.3 Modeling and system identification

Our dynamics model of the airplane is based on the model described in [Sobolic, 2009] ([Stevens and Lewis, 1992] is also a good reference for modeling fixed-wing airplanes). The model has 12 states:

$$\mathbf{x} = [x, y, z, \phi, \theta, \psi, \dot{x}, \dot{y}, \dot{z}, P, Q, R].$$

Here, $+x$ is in the forward direction, $+y$ is to the right and $+z$ is downwards as depicted in Figure 16(b) (this is the standard North-East-Down coordinate frame used in aeronautics). The states ϕ, θ, ψ are the roll, pitch and yaw angles. The variables P, Q, R are the components of the angular velocity expressed in the body coordinate frame. The control inputs of the model are the angles of the ailerons, rudder and elevator, along with the speed of the propellor. Thus, we have 4 control inputs since the ailerons are coupled to deflect in opposite directions by the commanded magnitude.

The airplane is treated as a rigid body with aerodynamic and gravitational forces acting on it. Here, we provide a brief description of our model of the aerodynamic forces:

Propellor thrust The thrust from the propellor is proportional to the square of the propellor speed. The constant of proportionality was obtained by hanging the airplane (with the propellor pointing downwards) from a digital fish scale and measuring the thrust produced for a number of different throttle settings.

Lift/drag on aerodynamic surfaces The lift and drag forces on the ailerons, rudder, elevator and tail of the airplane were computed using the flat-plate model. The flat-plate model was also used as a baseline for the lift and drag coefficients of the wings, but an angle-of-attack dependent correction term was added. This correction term was fit from experimental data obtained from passive (i.e., unactuated) flights in a manner similar to [Moore, 2014]. Since lift and drag forces are dependent on the airspeed over the aerodynamic surface, we need to take into account the effect of “propwash” (i.e. the airflow from the propellor). The relationships between the throttle speed command and the propellor downwash speed over the different control surfaces were measured using a digital anemometer in a manner similar to [Sobolic, 2009].

Body drag The drag on the airplane body is approximated as a quadratic drag term whose drag coefficient is fit from data.


As described above, many of the parameters in the model were obtained directly from physical experiments and measurements. However, some model parameters are more difficult to measure directly. These include the moments of inertia of the airplane and the coefficient of drag associated with the airplane body. The prediction-error minimization method in the Matlab System Identification Toolbox [Ljung, 2007] was used to fit these parameters and to fine-tune the measured parameters. In particular, we collected data from 15-20 flights (each lasting approximately 0.5–0.7 seconds) where the control inputs were excited using sinusoidal signals of varying frequency and amplitude.

8.4 Funnel validation

The first major goal of our hardware experiments was to demonstrate that our model of the SBach airplane is accurate enough to compute funnels that are truly meaningful for the hardware system.



(a) A depiction of the funnel that was validated on hardware. The funnel has been projected down to the $x - y - z$ coordinates of the state space and then re-projected onto the camera image.



(b) The value of the Lyapunov function (V) plotted as a function of time for 30 different trials of the airplane started from different initial conditions in the inlet of the funnel. The 1-sublevel set of the Lyapunov function corresponds to the funnel (i.e., a Lyapunov function of 1 or less corresponds to the airplane being inside the funnel). All 30 of the trajectories remain inside the computed funnel for the entire duration of the maneuver.

Figure 17: Validating funnels on the fixed-wing airplane.

To this end, we computed a funnel (shown in Figure 17(a)) for the airplane using the approach described in Section 4. We first estimated the set of initial states that the launcher mechanism causes the airplane to start off in (here, by “initial state” we mean the state of the airplane as soon as it has cleared the launcher mechanism). This was done by fitting an ellipsoid around the initial states observed from approximately 50 experimental trials. We then used direct collocation trajectory optimization [Betts, 2001] to design an open-loop maneuver that makes the airplane bend towards the left. The initial state of the trajectory is constrained to be equal to the mean of the experimentally observed initial states and the control inputs are constrained to satisfy the limits imposed by the hardware. Next we computed a time-varying LQR (TVLQR) controller around this nominal trajectory. This controller was tuned largely in simulation to ensure good tracking of the trajectory from the estimated initial condition set. The resulting closed-loop dynamics were then Taylor expanded around the nominal trajectory to degree 3 in order to obtain polynomial dynamics. Finally, we used SOS programming to compute the funnel depicted in Figure 17(a) using the time-sampled approximation described in Section 4.2 (with 10 time samples). The inlet of the funnel was constrained to contain the experimentally estimated set of initial conditions. We observed that the tuned TVLQR controller does not saturate the control inputs for the most part and thus we did not find it necessary to take input saturations into account in our funnel computation. Further, we wanted to assess the validity of our funnel for the nominal dynamics model of the airplane and did not take into account any uncertainty in the model. The funnel computation took approximately 1 hour.

We validate the funnel shown in Figure 17(a) with 30 experimental trials of the airplane ex-




Figure 18: Still images of the SBach flying through a funnel. The funnel has been projected down to the $x - y - z$ coordinates of the state space and then reprojected onto the camera image.

ecuting the maneuver corresponding to the funnel. The airplane is started off in different initial conditions in the inlet of the funnel and the TVLQR controller is applied for the duration of the maneuver. Figure 18 shows still images of a sample flight of the airplane executing the maneuver with the funnel superimposed onto the images. A video with a visualization of the funnel and flights through it is available online at <https://youtu.be/cESFpLgSb50>.

Figure 17(b) provides a more quantitative perspective on the flights. In particular, the figure shows the value of the Lyapunov function $V(\bar{x}, t)$ as a function of time achieved during the 30 experimental trials. Here, the Lyapunov function has been normalized so that the 1-level set corresponds to the boundary of the funnel. As the plot illustrates, all 30 trajectories remain inside the computed 12 dimensional funnel for the entire duration of the maneuver. This suggests that our model of the airplane is accurate enough to produce funnels that are indeed valid for the hardware system.

8.5 Obstacle avoidance experiments

The second major goal of our hardware experiments was to demonstrate the funnel library based real-time planning algorithm proposed in Section 6 on the obstacle avoidance task described in

Section 8.2. Our first step was to design a rich trajectory library consisting of a large number of different maneuvers. We initialized the library with the maneuver from Section 8.4 and augmented it by computing trajectories with varying final states. In particular, the library consists of 40 trajectories which were obtained by discretizing the final state in the y and z coordinates and using direct collocation trajectory optimization to compute locally optimal trajectories for the airplane model described in Section 8.3. The $x - y - z$ components of this trajectory library are depicted in Figure 19, with the trajectory from Section 8.4 highlighted in blue.




Figure 19: Trajectory library for the SBach airplane. The library consists of 40 different maneuvers. The maneuver from Section 8.4 is highlighted in blue.

For each trajectory in our library, we computed a TVLQR controller with the same state/action costs as the controller for the maneuver in Section 8.4. As in Section 8.4, we used SOS programming to compute funnels for each trajectory in the library. We note that since the dynamics of the airplane are symmetric with respect to reflection about the x axis, we were able to halve the amount of computation involved in constructing the trajectory/funnel library by exploiting this symmetry.

As mentioned in Section 8.2, the positions and geometry of the obstacles are not reported to the planner until the airplane has cleared the launcher. This forces planning decisions to be made in real-time. We use the planning algorithm described in Section 6 to choose a funnel from our library. The planner employs the QCQP-based algorithm from Section 6.1 for shifting funnels in the $x - y - z$ directions. As in the other examples considered in this paper, we use the FORCES Pro solver [Domahidi and Jerez, 2014] for our QCQP problems. If no satisfactory funnel is found by the planner, we revert to a “failsafe” option which involves switching off the propellor and gliding to a halt. A more sophisticated failsafe would be to attempt to transition to a “propellor-hang” mode. However, we found that our failsafe provides adequate protection to the airplane’s hardware as it usually glides on to the safety net before colliding with any obstacles. Finally, we note that since the experimental arena is quite limited in space, we do not replan funnels once one has been chosen; the airplane executes the feedback controller corresponding to the chosen funnel for the whole duration of the flight.

We tested our approach on 15 different obstacle environments of varying difficulty. These

environments are shown in Figures 20 and 21. The obstacles include poles of different lengths in varying orientations and also a hoop of diameter equal to 0.9 m (as a reference the airplane’s wingspan is 0.44 m , thus only leaving 0.23 m of margin on either side of the airplane assuming it passes through the center of the hoop). We model the poles as cuboids with heights equal to that of the poles and widths equal to the diameter. The hoop is approximated with eight polytopic segments (see Figure 23).

Figure 22 presents a more quantitative perspective on the obstacle environments. Here we have plotted a histogram of the gaps between obstacles in the environment and compared it to the wingspan of the airplane. In particular, for each obstacle in a given environment we consider the distance⁴ to the closest obstacle that is at least 5 cm away (the 5 cm threshold is chosen to prevent obstacles that are right next to each other being considered as having a small gap. For example, the obstacle closest to one of the horizontal poles in the first environment in Figure 20 should be the other horizontal pole and not the adjacent vertical pole). As the histogram illustrates, a significant fraction (approximately 35%) of the gaps are less than the airplane wingspan and about 66% of the gaps are less than two wingspans. Of course, it is worth noting that not all of the gaps greater than the wingspan are in fact negotiable (e.g., obstacles placed far apart along the x direction).

A video of the airplane traversing a few representative environments is available online at <https://youtu.be/cESFpLgSb50>. Out of the 15 environments the airplane was able to successfully negotiate 14 of them, thus demonstrating the efficacy of our real-time planner on this challenging task.

Figure 23 presents the output of our real-time planner on four of the more challenging environments. In particular we plot the funnel chosen by the planner (which have been shifted using the QCQP-based algorithm presented in Section 6.1) alongside an image sequence showing the airplane executing the plan. We note that the increased expressivity afforded to us by the ability to shift funnels in the cyclic coordinates has a large impact on the planner being able to find collision free funnels. Perhaps the best example of this is the environment which contains the hoop (top row of Figure 23). Without the ability to make small adjustments to the funnels, the chances of finding a collision free funnel are extremely low (we would require a funnel that passes almost exactly through the center of the hoop). Figure 24 compares the output of the planner with and without applying the QCQP-based algorithm for shifting funnels. As we can observe, the best funnel found in the former case is well in collision with the obstacles while in the latter case the planner is able to find a collision free funnel. This illustrates the crucial role that exploiting invariances plays in the success of the planner on this task.

As mentioned before, the airplane was able to successfully negotiate 14 out of our 15 environments. The single failure case occurred on the environment shown at the bottom of Figure 21, where the airplane clipped one of the poles close to the end of its flight. This failure can be attributed to the fact that the planner chose to execute one of a small handful of maneuvers that are more aggressive than the funnel we validated in Section 8.4. The controller saturates the control inputs significantly on this maneuver and hence violates our assumption about input limits not being reached. Hence this funnel is not entirely valid on the hardware system. While we could have taken actuator saturations into account while computing the funnels (see Section 4.3.3), we chose not to do so in order to reduce the computation time. In hindsight, a more careful treatment of the system would have included saturations when computing funnels.

⁴Here, the distance between obstacles is defined in the usual way for sets. In particular, for any pair of obstacles O_1 and O_2 , we define the distance between them as $\min_{x_1 \in O_1, x_2 \in O_2} \|x_1 - x_2\|_2$.



Figure 20: Environments 1-8 on which the online planning algorithm was tested. The obstacles include poles of different lengths in varying orientations and also a hoop of diameter equal to 0.9 m.




Figure 21: Environments 9-15 on which we tested our planning algorithm. The bottom most image shows the only failure case. Here the airplane brushed one of the obstacles on its way across the room.




Figure 22: Histogram of gaps between objects in our test environments compared with the airplane’s wingspan (0.44 m). As the plot shows, a significant number of the gaps between obstacles are less than then wingspan.

Finally, we note that on a small number of occasions during our experiments we observed the “failsafe” option being employed by the planner (i.e., switching off the throttle and gliding to a halt). This was typically caused by failures in the launching mechanism such as the airplane making contact with the operator’s hand on its way out of the launcher. Due to the altered initial conditions in these cases the airplane is not in the inlet of most or all of the funnels and thus has to resort to the failsafe. On these occasions we simply repeated the experiment to obtain a successful flight.

8.6 Implementation details

We end this section by mentioning a few implementation details that were important in achieving the results presented above. First, while the Vicon motion tracking system provides accurate position and orientation estimates at 120 Hz, a finite difference of these measurements can lead to noisy estimates of the derivative states. For our experiments we filtered the finite differences using a simple Luenberger observer [Luenberger, 1971].

Second, we observed a delay of approximately 55 ms in our closed-loop hardware system. While there are several ways to explicitly take this into account by adding delay states to our airplane model, we accommodate for the delay during execution by simulating our model forwards by 55 ms from the estimated current state and using this simulated future state to compute the current control input. This simple strategy is a common one and has previously been found to be effective in a wide range of applications [Horiuchi et al., 1999, Moore et al., 2014, Sipahi et al., 2012].









Figure 23: This figure depicts the planned funnel along with the polygonal obstacle representations for four of our fifteen test environments. Note that the funnels have been inflated to take into account the collision geometry of the airplane (modeled as a sphere of diameter equal to the wingspan).




Figure 24: The QCQP-based algorithm for shifting funnels (Section 6.1) plays a crucial role in the planner’s ability to find collision free funnels. Here we compare the output of the planner with and without applying the QCQP-based algorithm for shifting funnels. The best funnel found in the former case is well in collision with the obstacles while in the latter case the planner is able to find a collision free funnel. Note that the funnels chosen in the different cases correspond to different maneuvers.

9 Discussion and Conclusion

In this paper we have presented an approach for real-time motion planning in a priori unknown environments with dynamic uncertainty in the form of bounded parametric model uncertainty and external disturbances. The method augments the traditional trajectory library approach by constructing stabilizing controllers around the nominal trajectories in a library and computing outer approximations of reachable sets (funnels) for the resulting closed-loop controllers via sums-of-squares (SOS) programming. The pre-computed funnel library is then used to plan online by sequentially composing them together in a manner that ensures obstacles are avoided.

We have demonstrated our approach using extensive simulation experiments on a ground vehicle model. These experiments demonstrate that our approach can afford significant advantages over a trajectory-based approach. We also applied our approach to a quadrotor model and demonstrated how for certain classes of environments we can guarantee that the system will fly forever in a collision-free manner. We have also validated our approach using thorough hardware experiments on a small fixed-wing airplane flying through previously unseen cluttered environments at high speeds. To our knowledge, the resulting hardware demonstrations on a fixed-wing airplane constitute one of the first examples of provably safe and robust control for robotic systems with complex nonlinear dynamics that need to plan in realtime in environments with complex geometric constraints. It is also worth noting that while one often associates robustness with very conservative behavior, our hardware experiments demonstrate that this need not be the case. In particular, the airplane performs some very aggressive maneuvers while still being robust.

9.1 Challenges and extensions

9.1.1 Numerical difficulties

There are a number of challenges associated with our approach. One of the main difficulties in implementing our method is the relative immaturity of solvers for semidefinite programs (SDPs). While recently released software such as the MOSEK solver [Mosek, 2010] have improved the speed

with which solutions can be obtained, SDP solvers are still in their infancy as compared to solvers for Linear Programs (LPs). Thus, numerical issues (e.g., due to the scaling of problem data) inevitably arise in practice and must be dealt with in a relatively ad hoc manner (e.g., rescaling problem data or removing redundant decision variables). However, preprocessing of sums-of-squares (SOS) programs and SDPs is an active area of inquiry [Permenter and Parrilo, , Lofberg, 2009, Seiler et al., 2013] and solver technology is bound to improve as SDPs are more widely adopted in practice.

9.1.2 Robots with complex geometries

Another challenge associated with our method has to do with implementing the approach on robots with complex kinematics/geometry. One option (as described in Section 6) is to project the funnel onto the configuration space (C-space) of the system and perform collision checking against the C-space representation of obstacles. However, computing C-space representations of obstacles is typically challenging for non-trivial geometries, especially since it needs to be done as obstacles are reported in real-time.

For the examples considered in this paper, the geometry of the robots were approximated relatively accurately by spheres. This allowed us to project the funnel onto the $x - y - z$ space (in contrast to the full configuration space of the system) and inflate this projection by the radius of the corresponding sphere. These inflated funnels were then used for collision-checking during real-time planning. For robots with complex geometries, inflating the funnel in $x - y - z$ space in this manner is not an entirely straightforward operation. However, this is potentially a more promising approach than performing collision-checking against C-space obstacles since the inflation can be computed in the offline computation stage as it depends only on the geometry of the robot and not the obstacles.

9.1.3 Designing funnel libraries

There is an inherent tradeoff in our approach between the richness of the funnel library and the amount of computation that needs to be performed in real-time in order to be able to search through it. For extremely large funnel libraries, it may be computationally difficult to search through all the funnels while planning online. It is thus important to explore ways in which one could speed up this search. For example, one could exploit the observation that funnels in close proximity to a funnel that is in collision are also likely to be in collision and use this to choose the sequence in which to evaluate funnels for collisions. Further, given certain features of the environment one may be able to predict which funnels are likely to be in collision (without actually performing collision-checking) and evaluate these funnels first. These intuitions have been formalized and exploited using the theory of submodular optimization in the context of trajectory libraries [Dey et al., 2012, Dey, 2015]. The approach allows one to optimize the sequence in which trajectories are evaluated and should be generalizable to funnel libraries as well.

A closely related question is how to choose the funnels in the library in the first place. As we observed in Section 7.2 for the quadrotor system, one can derive relatively simple geometric conditions on the environment in order for us to be able to guarantee that the system will be able to navigate through it without collisions. If we know a priori that our environment will satisfy these geometric conditions, this provides a way to check if our funnel library is sufficiently rich. However, for real-world environments (e.g., forests) we may not be able to make such assumptions. Instead, we might have a *generative* (probabilistic) model of our environments and could use this

to evaluate/design our funnel library. For instance, it is known that the locations of trees in a forest are modeled well by spatial Poisson processes [Stoyan and Penttinen, 2000]. In such settings, it may be possible to design a randomized algorithm for generating the funnel library where one samples different realizations of the environment, searches through the existing library to find a collision-free funnel for the sampled environment, and adds a new funnel to the library when such a funnel doesn't exist.

9.1.4 Probabilistic guarantees

Throughout this paper, we have assumed that all disturbances and uncertainty are bounded with probability one. In practice, this assumption may either not be fully valid or could lead to overly conservative performance. In such situations, it is more natural to provide guarantees of a probabilistic nature. Recently, results from classical super-martingale theory have been combined with sums-of-squares programming in order to compute such probabilistic certificates of finite time invariance [Steinhardt and Tedrake, 2012], i.e. provide upper bounds on the probability that a stochastic nonlinear system will leave a given region of state space. Combining the techniques presented in [Steinhardt and Tedrake, 2012] with the approach presented in this work to perform robust online planning on stochastic systems will be the subject of future work.

9.1.5 Reasoning about perception systems

In this paper, we have focused most of our attention on reasoning about the real-time planning and control systems. In particular, we assumed that the robot is equipped with a perception system that reports (at runtime) regions in which obstacles are guaranteed to lie within. This may not always be a valid assumption; in practice, perceptual systems can often have false negatives and fail to report an obstacle in the environment. In such scenarios, it is unreasonable to expect to guarantee with probability one that the system will remain collision-free and we can only hope for probabilistic guarantees. One could envision modeling this perceptual uncertainty by considering an occupancy map with probabilities of occupancy associated with each voxel. Modeling and reasoning about such perceptual uncertainty remains an important open problem.

9.1.6 Real-time computation of funnels

Due to the computation time associated with funnels, the approach presented in this work had two phases: an offline phase for computing funnels and an online stage for real-time planning with funnels. Recently, more scalable alternatives to SOS programming have been introduced [Ahmadi and Majumdar, 2016, Ahmadi and Majumdar, 2014]. These alternatives rely on *linear* and *second-order cone programming* instead of semidefinite programming. This makes it possible to obtain large computational gains in terms of scalability and running time using DSOS and SDSOS programming as compared to SOS programming (see [Majumdar et al., 2014] for comparisons of DSOS/SDSOS with SOS on large-scale controls applications). This raises the interesting possibility of being able to dispense with the offline computation stage and compute a funnel and feedback controller *in real-time* based on the geometry of the environment. Currently, this would only be feasible for relatively low-dimensional systems but may be an exciting direction to pursue.

We believe that the work presented in this paper has the potential to be deployed on real robots to make them operate safely in real-world environments. Our hope is that by building upon this work and pursuing the directions for future research presented above we can make this a reality.

Acknowledgements

The authors are grateful to Tomas Lozano-Perez, Emilio Frazzoli, Claire Tomlin, and the members of the Robot Locomotion Group at MIT for a number of helpful discussions on this work. The authors would also like to thank Michael Posa for help formulating the QCQP for shifting funnels presented in Section 6.1; Adam Bry and Charlie Richter for help setting up the USB RC transmitter used in the experiments described in Section 8; and Andy Barry for advice on choosing the fixed-wing airplane used in Section 8 and his help with the launching mechanism for the airplane.

Funding

This work was supported by the Office of Naval Research [MURI grant number N00014-10-1-095].

References

- [Agha-Mohammadi et al., 2014] Agha-Mohammadi, A.-A., Chakravorty, S., and Amato, N. (2014). FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *The International Journal of Robotics Research*, 33(2):268–304.
- [Ahmadi and Majumdar, 2014] Ahmadi, A. A. and Majumdar, A. (2014). DSOS and SDSOS optimization: LP and SOCP-based alternatives to sum of squares optimization. In *48th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–5. IEEE.
- [Ahmadi and Majumdar, 2016] Ahmadi, A. A. and Majumdar, A. (2016). DSOS and SD-SOS optimization: More tractable alternatives to SOS optimization. *In preparation* (<http://aaa.lids.mit.edu/publications>).
- [Althoff et al., 2015] Althoff, D., Althoff, M., and Scherer, S. (2015). Online safety verification of trajectories for unmanned flight with offline computed robust invariant sets. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 3470–3477. IEEE.
- [Althoff and Dolan, 2014] Althoff, M. and Dolan, J. (2014). Online verification of automated road vehicles using reachability analysis. *IEEE Transactions on Robotics*, 30(4):903–918.
- [Althoff et al., 2008] Althoff, M., Stursberg, O., and Buss, M. (2008). Reachability analysis of non-linear systems with uncertain parameters using conservative linearization. In *IEEE Conference on Decision and Control (CDC)*, pages 4042–4048. IEEE.
- [Barry, 2016] Barry, A. J. (2016). *High-Speed Autonomous Obstacle Avoidance with Pushbroom Stereo*. PhD thesis, Massachusetts Institute of Technology.
- [Barry et al., 2012] Barry, A. J., Majumdar, A., and Tedrake, R. (2012). Safety verification of reactive controllers for UAV flight in cluttered environments using barrier certificates. In *Proceedings of the 2012 IEEE International Conference on Robotics and Automation (ICRA)*.

[Ben-Tal and Nemirovski, 2001] Ben-Tal, A. and Nemirovski, A. (2001). *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, volume 2. SIAM.

[Berenson et al., 2007] Berenson, D., Diankov, R., Nishiwaki, K., Kagami, S., and Kuffner, J. (2007). Grasp planning in complex scenes. In *International Conference on Humanoid Robots*, pages 42–48. IEEE.

[Berg et al., 2012] Berg, J. V. D., Patil, S., and Alterovitz, R. (2012). Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research*, 31(11):1263–1278.

[Betts, 2001] Betts, J. T. (2001). *Practical Methods for Optimal Control Using Nonlinear Programming*. SIAM Advances in Design and Control. Society for Industrial and Applied Mathematics.

[Blackmore et al., 2006] Blackmore, L., Li, H., and Williams, B. (2006). A probabilistic approach to optimal robust path planning with obstacles. In *American Control Conference, 2006*, page 7 pp.

[Blekherman et al., 2013] Blekherman, G., Parrilo, P. A., and Thomas, R. R. (2013). *Semidefinite optimization and convex algebraic geometry*, volume 13. SIAM.

[Boyd and Vandenberghe, 2004] Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.

[Brooks, 1982] Brooks, R. (1982). Symbolic error analysis and robot planning. *The International Journal of Robotics Research*, 1(4):29–78.

[Brown, 2003] Brown, C. W. (2003). QEPCAD B: A program for computing with semi-algebraic sets using CADs. *ACM SIGSAM Bulletin*, 37(4):97–108.

[Bry and Roy, 2011] Bry, A. and Roy, N. (2011). Rapidly-exploring random belief trees for motion planning under uncertainty. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Shanghai, China.

[Burridge et al., 1999] Burridge, R. R., Rizzi, A. A., and Koditschek, D. E. (1999). Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research*, 18(6):534–555.

[Charnes and Cooper, 1959] Charnes, A. and Cooper, W. W. (1959). Chance-constrained programming. *Management Science*, 6(1):pp. 73–79.

[Chen and Tomlin, 2015] Chen, M. and Tomlin, C. (2015). Exact and efficient hamilton-jacobi reachability for decoupled systems. *arXiv preprint arXiv:1503.05933*.

[Collins and Hong, 1998] Collins, G. E. and Hong, H. (1998). *Partial cylindrical algebraic decomposition for quantifier elimination*. Springer.

[Coumans, 2014] Coumans, E. (2014). Bullet physics library. *bulletphysics. org*, 4(6).

[Dey, 2015] Dey, D. (2015). *Predicting Sets and Lists: Theory and Practice*. PhD thesis, Carnegie Mellon University.

[Dey et al., 2011] Dey, D., Liu, T., Sofman, B., and Bagnell, D. (2011). Efficient optimization of control libraries. Technical report, Technical Report (CMU-RI-TR-11-20).

[Dey et al., 2012] Dey, D., Liu, T. Y., Hebert, M., and Bagnell, A. J. (2012). Contextual sequence prediction with application to control library optimization. In Proceedings of Robotics Science and Systems.

[Dey et al., 2014] Dey, D., Shankar, K. S., Zeng, S., Mehta, R., Agcayazi, M. T., Eriksen, C., Daftary, S., Hebert, M., and Bagnell, J. A. (2014). Vision and learning for deliberative monocular cluttered flight. *arXiv preprint arXiv:1411.6326*.

[Domahidi and Jerez, 2014] Domahidi, A. and Jerez, J. (2014). FORCES Professional. Embotech GmbH (<http://embotech.com/FORCES-Pro>).

[Dubins, 1957] Dubins, L. (1957). On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):497–516.

[Frazzoli et al., 2005] Frazzoli, E., Dahleh, M., and Feron, E. (2005). Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE Transactions on Robotics*, 21(6):1077–1091.

[Frazzoli et al., 2002] Frazzoli, E., Dahleh, M. A., and Feron, E. (2002). Real-Time Motion Planning for Agile Autonomous Vehicles. *Journal of Guidance, Control, and Dynamics*, 25(1):116–129.

[Gao et al., 2013] Gao, S., Kong, S., and Clarke, E. M. (2013). dReal: An SMT solver for nonlinear theories over the reals. In *Automated Deduction–CADE-24*, pages 208–214. Springer.

[Gillula et al., 2010] Gillula, J., Huang, H., Vitus, M., and Tomlin, C. (2010). Design of guaranteed safe maneuvers using reachable sets: Autonomous quadrotor aerobatics in theory and practice. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1649–1654. IEEE.

[Girard, 2005] Girard, A. (2005). Reachability of uncertain linear systems using zonotopes. In *Hybrid Systems: Computation and Control*, pages 291–305. Springer.

[Green and Kelly, 2007] Green, C. and Kelly, A. (2007). Toward optimal sampling in the space of paths. In *13th International Symposium of Robotics Research*.

[Guibas et al., 2009] Guibas, L., Hsu, D., Kurniawati, H., and Rehman, E. (2009). Bounded uncertainty roadmaps for path planning. In *Algorithmic Foundation of Robotics VIII*, pages 199–215. Springer.

[Horiuchi et al., 1999] Horiuchi, T., Inoue, M., Konno, T., and Namita, Y. (1999). Real-time hybrid experimental system with actuator delay compensation and its application to a piping system with energy absorber. *Earthquake Engineering & Structural Dynamics*, 28(10):1121–1141.

[Jacobs and Canny, 1990] Jacobs, P. and Canny, J. (1990). Robust motion planning for mobile robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2–7. IEEE.

[Kaelbling et al., 1998] Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101.

[Karaman and Frazzoli, 2011] Karaman, S. and Frazzoli, E. (2011). Sampling-based algorithms for optimal motion planning. *Int. Journal of Robotics Research*, 30:846–894.

[Kobilarov, 2012] Kobilarov, M. (2012). Cross-entropy motion planning. *The International Journal of Robotics Research*, 31(7):855–871.

[Konidaris et al., 2011] Konidaris, G., Kuindersma, S., Grupen, R., and Barto, A. (2011). Robot learning from demonstration by constructing skill trees. *The International Journal of Robotics Research*.

[Kuffner and Lavalle, 2000] Kuffner, J. and Lavalle, S. (2000). RRT-connect: An efficient approach to single-query path planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 995–1001.

[Kurzhanski and Varaiya, 2000] Kurzhanski, A. and Varaiya, P. (2000). Ellipsoidal techniques for reachability analysis: internal approximation. *Systems & control letters*, 41(3):201–211.

[Kuwata et al., 2008] Kuwata, Y., Fiore, G. A., Teo, J., Frazzoli, E., and How, J. P. (2008). Motion planning for urban driving using RRT. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1681–1686. IEEE.

[Latombe et al., 1991] Latombe, J.-C., Lazanas, A., and Shekhar, S. (1991). Robot motion planning with uncertainty in control and sensing. *Artificial Intelligence*, 52(1):1–47.

[Li and Williams, 2008] Li, H. X. and Williams, B. (2008). Generative planning for hybrid systems based on flow tubes. In *International Conference on Automated Planning and Scheduling (ICAPS)*, pages 206–213.

[Liu and Atkeson, 2009] Liu, C. and Atkeson, C. (2009). Standing balance control using a trajectory library. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 3031–3036. IEEE.

[Ljung, 2007] Ljung, L. (2007). System identification toolbox for use with matlab.

[Lofberg, 2009] Lofberg, J. (2009). Pre- and post-processing sum-of-squares programs in practice. *IEEE Transactions On Automatic Control*, 54(5):1007–.

[Lozano-Perez et al., 1984] Lozano-Perez, T., Mason, M., and Taylor, R. (1984). Automatic synthesis of fine-motion strategies for robots. *The International Journal of Robotics Research*, 3(1):3–24.

[Luenberger, 1971] Luenberger, D. (1971). An introduction to observers. *IEEE Transactions on Automatic Control*, 16(6):596–602.

[Majumdar et al., 2013] Majumdar, A., Ahmadi, A. A., and Tedrake, R. (2013). Control design along trajectories with sums of squares programming. In *Proceedings of the 2013 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4054–4061.

[Majumdar et al., 2014] Majumdar, A., Ahmadi, A. A., and Tedrake, R. (2014). Control and verification of high-dimensional systems with DSOS and SDSOS programming. In *Proceedings of the 53rd Conference on Decision and Control (CDC)*.

[Majumdar and Tedrake, 2012] Majumdar, A. and Tedrake, R. (2012). Robust online motion planning with regions of finite time invariance. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, page 16, Cambridge, MA.

[Malone et al., 2013] Malone, N., Manavi, K., Wood, J., and Tapia, L. (2013). Construction and use of roadmaps that incorporate workspace modeling errors. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 1264–1271. IEEE.

[Mason, 1985] Mason, M. (1985). The mechanics of manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 544–548. IEEE.

[Mayne et al., 2005] Mayne, D., Seron, M., and Rakovic, S. (2005). Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224.

[Mellinger et al., 2010] Mellinger, D., Michael, N., and Kumar, V. (2010). Trajectory generation and control for precise aggressive maneuvers with quadrotors. In *Proceedings of the 12th International Symposium on Experimental Robotics (ISER 2010)*.

[Missiuro and Roy, 2006] Missiuro, P. and Roy, N. (2006). Adapting probabilistic roadmaps to handle uncertain maps. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1261–1267. IEEE.

[Mitchell et al., 2005] Mitchell, I., Bayen, A., and Tomlin, C. (2005). A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957.

[Moore, 2014] Moore, J. (2014). *Robust Post-Stall Perching with a Fixed-Wing UAV*. PhD thesis, Massachusetts Institute of Technology.

[Moore et al., 2014] Moore, J., Cory, R., and Tedrake, R. (2014). Robust post-stall perching with a simple fixed-wing glider using lqr-trees. *Bioinspiration and Biomimetics*, 9(2):15.

[Mosek, 2010] Mosek, A. (2010). The mosek optimization software. *Online at <http://www.mosek.com>*, 54.

[Ny and Pappas, 2012] Ny, J. L. and Pappas, G. (2012). Sequential composition of robust controller specifications. In *International Conference on Robotics and Automation (ICRA)*. IEEE.

[Oldewurtel et al., 2008] Oldewurtel, F., Jones, C. N., and Morari, M. (2008). A tractable approximation of chance constrained stochastic mpc based on affine disturbance feedback. In *IEEE Conference on Decision and Control*, pages 4731–4736. IEEE.

[Parrilo, 2000] Parrilo, P. A. (2000). *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology.

[Patil et al., 2015] Patil, S., Kahn, G., Laskey, M., Schulman, J., Goldberg, K., and Abbeel, P. (2015). Scaling up gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation. In *Algorithmic Foundations of Robotics XI*, pages 515–533. Springer.

[Permenter and Parrilo,] Permenter, F. and Parrilo, P. Basis selection for SOS programs via facial reduction and polyhedral approximations. *Proceedings of the IEEE Conference on Decision and Control, December 2014*.

[Platt et al., 2012] Platt, R., Kaelbling, L., Lozano-Perez, T., and Tedrake, R. (2012). Non-gaussian belief space planning: Correctness and complexity. In *Proceedings of the 2012 IEEE International Conference on Robotics and Automation (ICRA)*.

[Platt et al., 2010] Platt, R., Tedrake, R., Kaelbling, L., and Lozano-Perez, T. (2010). Belief space planning assuming maximum likelihood observations. In *Proceedings of Robotics: Science and Systems*.

[Prajna, 2006] Prajna, S. (2006). Barrier certificates for nonlinear model validation. *Automatica*, 42(1):117 – 126.

[Prentice and Roy, 2007] Prentice, S. and Roy, N. (2007). The Belief Roadmap: Efficient planning in linear POMDPs by factoring the covariance. In *Proceedings of the 12th International Symposium of Robotics Research (ISRR)*.

[Satzinger et al., 2016] Satzinger, B. W., Lau, C., Byl, M., and Byl, K. (2016). Experimental results for dexterous quadruped locomotion planning with RoboSimian. In *Experimental Robotics*, pages 33–46. Springer.

[Seiler et al., 2013] Seiler, P., Zheng, Q., and Balas, G. (2013). Simplification methods for sum-of-squares programs. *arXiv preprint arXiv:1303.0714*.

[Sermanet et al., 2008] Sermanet, P., Scoffier, M., Crudele, C., Muller, U., and LeCun, Y. (2008). Learning maneuver dictionaries for ground robot planning. In *Proc. 39th International Symposium on Robotics (ISR08)*.

[Shkolnik, 2010] Shkolnik, A. (2010). *Sample-Based Motion Planning in High-Dimensional and Differentially-Constrained Systems*. PhD thesis, MIT.

[Sipahi et al., 2012] Sipahi, R., Vyhlídal, T., Niculescu, S.-I., and Pepe, P. (2012). *Time Delay Systems: Methods, Applications and New Trends*, volume 423. Springer.

[Sobolic, 2009] Sobolic, F. M. (2009). Agile flight control techniques for a fixed-wing aircraft. Master’s thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, Cambridge MA.

[Steinhardt and Tedrake, 2012] Steinhardt, J. and Tedrake, R. (2012). Finite-time regional verification of stochastic nonlinear systems. *International Journal of Robotics Research*, 31(7):901–923.

[Stevens and Lewis, 1992] Stevens, B. and Lewis, F. (1992). *Aircraft Control and Simulation*. John Wiley & Sons, Inc.

[Stolle and Atkeson, 2006] Stolle, M. and Atkeson, C. (2006). Policies based on trajectory libraries. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*. IEEE.

[Stoyan and Penttinen, 2000] Stoyan, D. and Penttinen, A. (2000). Recent applications of point process methods in forestry statistics. *Statistical Science*, pages 61–78.

[Tedrake et al., 2010] Tedrake, R., Manchester, I. R., Tobenkin, M. M., and Roberts, J. W. (2010). LQR-Trees: Feedback motion planning via sums of squares verification. *International Journal of Robotics Research*, 29:1038–1052.

[Tobenkin et al., 2011] Tobenkin, M. M., Manchester, I. R., and Tedrake, R. (2011). Invariant funnels around trajectories using sum-of-squares programming. *Proceedings of the 18th IFAC World Congress, extended version available online: arXiv:1010.3013 [math.DS]*.

[Toit and Burdick, 2011] Toit, N. D. and Burdick, J. (2011). Probabilistic collision checking with chance constraints. *IEEE Transactions on Robotics*, 27(4):809–815.

[Vandenberghe and Boyd, 1996] Vandenberghe, L. and Boyd, S. (1996). Semidefinite programming. *SIAM Review*, 38(1):49–95.

[Vitus and Tomlin, 2011] Vitus, M. and Tomlin, C. (2011). Closed-loop belief space planning for linear, gaussian systems. In *International Conference on Robotics and Automation (ICRA)*, pages 2152–2159. IEEE.

[Vitus and Tomlin, 2012] Vitus, M. and Tomlin, C. (2012). A hybrid method for chance constrained control in uncertain environments. In *Conference on Decision and Control (CDC)*, pages 2177–2182. IEEE.

[Yazarel and Pappas, 2004] Yazarel, H. and Pappas, G. (2004). Geometric programming relaxations for linear system reachability. In *Proceedings of the American Control Conference*, volume 1, pages 553–559. IEEE.